ADVANCED RESEARCH PROJECTS AGENCY
3701 NORTH FAIRFAX DRIVE
ARLINGTON, VA 22203-1714

June 8, 1993

UNIVERSITY OF SOUTHERN CALIFORNIA
INFORMATION SCIENCES INSTITUTE
4676 Admiralty Way
Marina del Rey, CA 90292
Attn: Security

Dear Sir:

Reference is made to the DTIC-DDA letter dated February 19, 1993 concerning the following report(s) sent to DTIC without an approved distribution statement:

FINAL TR "DARPA/ISTO RAPID VLSI IMPLEMENTATION"
December 1991

The report listed above was cleared for open publication by the Directorate for Security Review, OASD(PA) on Mar 18, 1993, Case No. 93-S-0636. WITH DISCLAIMER.

All copies of this report should carry the following Distribution Statement (A) as follows:

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

A completed SF 298 should appear in each technical report produced under contract to ARPA.

Sincerely,

G. T. Winn
Technical Information Officer

Enclosures:
1. DTIC-DDA Letter
2. Report a/s

CC: DTIC

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE

3. REPORT TYPE AND DATES COVERED

4. TITLE AND SUBTITLE

5. FUNDING NUMBERS

6. AUTHOR(S)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

14. SUBJECT TERMS

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT

18. SECURITY CLASSIFICATION OF THIS PAGE

19. SECURITY CLASSIFICATION OF ABSTRACT

20. LIMITATION OF ABSTRACT

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to ***stay within the lines*** to meet ***optical scanning requirements***.

Block 1. Agency Use Only *(Leave blank)*.

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C	- Contract	**PR**	- Project
G	- Grant	**TA**	- Task
PE	- Program Element	**WU**	- Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. *(If known)*

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."
DOE - See authorities.
NASA - See Handbook NHB 2200.2.
NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.
DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.
NASA - Leave blank.
NTIS - Leave blank.

Block 13. Abstract. Include a brief *(Maximum 200 words)* factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code *(NTIS only)*.

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.

(12

UNIVERSITY OF SOUTHERN CALIFORNIA

INFORMATION SCIENCES INSTITUTE
4676 Admiralty Way Marina del Rey, CA 90292

Sponsored by Defense Advanced Projects Agency (DOD)
DARPA/CSTO

DARPA/ISTO RAPID VLSI IMPLEMENTATION

Final Technical Report

Date of Report: December 1991

Project: Advanced Production Technology
ARPA Order No: 6132 (PL)

Issued by: AFCMD/KCC
Under Contract Number: F29601-87-C-0069
Period of Performance: 09/28/87 - 09/30/91

Primary Investigator
Robert Parker
(310) 822-1511
Parker@venera.isi.edu

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

93-S-0636

(12

!NIVERSITY OF SOUTHERN CALIFORNIA

INFORMATION SCIENCES INSTITUTE

4676 Admiralty Way Marina del Rey, CA 90292

Sponsored by Defense Advanced Projects Agency (DOD)

DARPA/CSTO

DARPA/ISTO RAPID VLSI IMPLEMENTATION

Final Technical Report

Date of Report: December 1991

Project: Advanced Production Technology

ARPA Order No: 6132 (PL)

Issued by: AFCMD/KCC

Under Contract Number: F29601-87-C-0069

Period of Performance: 09/28/87 - 09/30/91

Primary Investigator

Robert Parker

(310) 822-1511

Parker@venera.isi.edu

93-S-0636

Final Report
Table of Contents

APT Project Goals

The overall goal of APT is to investigate and develop high-performance system packaging technologies and production approaches to support future increases in computer complexity and system clock rates and to accelerate the incorporation of these approaches in new architecture development.

Specific APT goals include:

- development and demonstration of a system design methodology that factors packaging into system design early in the design cycle rather than treating packaging as a post-design process.
- providing DARPA-sponsored architecture research teams with high-performance packaging technology by undertaking small-scale, closely coupled, development efforts that demonstrate methodologies for improved system performance. These technology demonstrations, called Collaborative Development Efforts (CDE), have the following goals:
  - minimizing the risk to collaborating partners by conducting parallel experiments.
  - maximizing future benefit to the collaborating partners through closely coupled, lock-step development.
  - validating and accelerating the incorporation of advanced packaging technology in new architectures through early experiments.
  - encouraging system architects to design new architectures which are optimized for advanced packaging techniques.
- establishment of a local advanced packaging technology base supplied from commercial sources. This technology base will provide characterization and models for new technologies, as well as procedures and sources for assembly and test of prototype and production systems.
- providing access to this technology base through the creation of a service called Packaging Feasibility Studies (PFS). The PFS will provide relevant packaging information to systems designers early in the design cycle, thus allowing more informed system design and partitioning decisions.
- support for this technology base by developing a flexible, multi-purpose, low-cost probe station environment suitable for testing demonstration system components, laser characterizing chips, and providing remote chip diagnostic support.

Packaging Feasibility Studies

INTRODUCTION

APT has devised the Packaging Feasibility Study (PFS) to formalize a directed response to application-specific requests. The PFS is conducted by asking designers for specifications of a particular system, including chip count, power, speed, interconnnect requirements, technology preference, and environmental constraints. The PFS is a report generated by applying the constraints of various packaging approaches onto the desired system specifications. The resulting document provides designers with trade-offs early in the design cycle, which allows the design process to proceed toward a final design that can be manufactured. The report covers required air flow for certain die temperatures; proposed packaging technologies for chips, boards, and cabinets; proposed die-attach methods; proposed design partitioning; proposed mechanical package design; and proposed interconnect technology. APT has conducted many PFSs for the DARPA community, mostly on multiprocessor architectures. The summary of several APT PFSs are described below.

AT&T

AT&T will use an internal hybrid process to achieve the desired packaging density for the ASPEN multiprocessor. This approach is similar to what APT would recommend that they do. An APT parallel cooperative packaging effort using the same die but a commercial hybrid approach would lower the risk to DARPA by providing an *alternative source for* hybrids. It is not clear that performance would be improved by assuming more packaging risk because the DSP32C (the ASPEN processor chip) performance appears to be the determining factor.

INTEL/CMU

Intel recognized that higher performance packaging approaches for their IWARP systolic machine was necessary. Packaging issues ranging from die-attach methods to system packaging were discussed.

ENCORE COMPUTER - LYNX

A draft study has been completed for a packaging approach proposed for the Encore LYNX. In addition, a thermal mock-up has been built and is currently undergoing evaluation.

The results of this study contain proprietary information furnished under one or more non-disclosure agreements.

AMETEK/CALTECH

CANTER

Caltech has developed families of multiprocessor cube interconnect architectures with varying degrees of processing node complexity. One particular design, the Canter Engine, in-

volved a custom VLSI routing chip (FMRC), a custom RISC processor, a custom list processor (LT-1), and several commercial RAM chips per node. The total device count of approximately 15 chips and wide data paths makes this architecture very interesting as a packaging application. APT examined use of VLSI designed at Caltech coupled with fast SRAMs to build a one-dimensional cube. This approach was intended to produce a 160-MIPS computing engine in a package 1"x2"x2."

This simple node provides very high packing-factor and high performance. The application of a few-chips-per-node is a near-ideal hybrid application, requiring mixed technology that is too large for wafer-scale approaches. The hybrid approach can produce powerful multi-node workstation-size machines with little or no software development.

MOSAIC

An early prototype of the MOSAIC multiprocessor cube developed at Caltech used SIP (single in-line package) memory technology. Since the machine nodes were dominated by physical memory, APT would have achieved perhaps 3/4 of the density of a full hybrid approach. The risks of a hybrid approach would not be warranted for this modest increase in packaging density.

UNIVERSITY OF TEXAS

A hybrid-based machine proposed by Bill Athas was "sized" by APT. The proposed machine was similar to Chuck Seitz's work at Caltech except for wider interconnect busses. The wider busses create higher bandwidth communications between nodes which is essential to match the increased performance of each node achieved through hybrid packaging. This application would be excellent for a hybrid approach.

BERKELEY

The hybrid-based Aquarius III Prolog machine proposed by Vason Srini was "sized". This machine was memory-intensive and required specialized packaging at several levels. APT proposed packaging that built stacked hybrid modules for each processor node and to used button technology for the interconnect busses.

TRW

TRW was interested in packaging a switch design. Data switch packaging is unique because the switching logic is typically minimal while the datapath I/O requirements are enormous. In the TRW switch, the VLSI devices require at least 500 I/O pins, meaning that multilevel TAB or flip-chip technology would be required. A very dense, high-performance switch could be built with CMOS logic and direct die attach methods.

MANUFACTURING SUPPORT FOR PACKAGING

Robert Parker presented a talk at the Center for Robotics at UC Santa Barbara. The audience was mostly mechanical engineers specializing in robotics or mechanical assembly.

Their interest was in understanding the problems associated with automated manufacture and robotic assembly of advanced packaging required to build next-generation computing machines. The talk outlined many of the mechanical problems facing the users of packaging technologies such as polyimide on silicon and button interconnect.

Exchange of ideas between materials and process researchers and manufacturing companies is essential to the development and acceptance of new packaging technologies. This particular presentation achieved three results. First, that particular audience was made aware of real problems in advanced packaging. Second, the meeting created the awareness that a white paper on the subject should be created. Third, a packaging assembly example used in the talk was apparently an excellent match for a particular robotics effort at UCSB.

PACKAGING WHITE PAPER

APT effort was directed toward continued evaluation of commercially available packaging technology. A report was written that proposes how technology at each packaging level should be accessed based on these evaluations. This report, submitted to DARPA as a "white paper" on packaging was a basis for continued refinement of capabilities required of a laboratory supporting advanced system prototyping. APT continued to define and develop the internal infrastructure required to evaluate and characterize advanced packaging technology.

GaAs TESTING

A package evaluation experiment was conducted in ISI's class 10,000 clean room facility. High-frequency GaAs test chips from the University of Utah were assembled in controlled-impedance packages manufactured by Triquint, and the parts tested using a matched high-frequency card mounted on the low-cost probe station environment. The experiment demonstrated the transmission of 45 megahertz signals through package pins into a glass-epoxy circuit board with remarkably little signal degradation. Signal risetimes of 500 picoseconds indicated that clock rates of 250 to 500 megahertz could be supported with this technology.

SARNOFF RESEARCH CENTER

The Sarnoff Research Center, under contract to DARPA, developed the Princeton Engine. Conventional technology had been used to package this machine. APT completed a Packaging Feasibility Study that presented short-term, alternative approaches that could be rapidly moved into production.

Two more aggressive approaches were discussed for future development of a multi-TeraOp machine: repackaging the existing design to achieve a system volume of one cubic foot and developing a new packaging strategy for a much higher performance machine.

Using the original design, existing memory modules would be replaced to achieve a factor-of-four improvement in machine memory capacity. This would be easily achieved with APT

technology and would be a low-cost, direct-plug-compatible enhancement, increasing the available memory to match the addressing capability of the existing machine.

The second APT proposal would use VLSI and MCM technology to shrink original design to the goal of one cubic foot. This complete system repackage was attractive to APT because of the already high level of integration. The same silicon would have been used for a high-density version.

Size reduction resulting from the repackage effort would allow the existing design to be inserted more easily into size-sensitive military systems. However, the overall performance of the machine would not have increased significantly as a result of repackaging because of performance of a multiplier buried inside one of the gate arrays.

The Princeton Machine architecture was ideal for advanced packaging because it was organized in "slices," with relatively few chips per slice and relatively few interconnect wires between slices. APT has found that systems with high levels of integration lend themselves to three-dimensional, high-density packaging.

UC SANTA BARBARA SHUNT

Originally, an MCM-based approach was proposed to support the 39 chips required for each switch node. Because of the small volumes projected for the initial fabrication run, wire-bonding had been proposed as the die interconnect method. Thermal analysis had indicated that a stacked MCM approach would acceptably accommodate both the thermal requirements and the high I/O counts. The MCM proposed originally was slightly larger than 2 inches and provided support for the over 1000 I/O interconnects. Costs of design and fabrication were provided and sources of technology were identified. A report was completed and forwarded to UCSB and to DARPA.

It developed that the cost of the MCM approach was beyond the fabrication budget established for the project, meaning that a lower-cost approach was required. APT accordingly undertook a second PFS to devise an approach that would meet the cost constraints and could be fabricated within the time remaining of the project implementation schedule.

The SHUNT Study evolved into a Collaborative Development Effort to implement the higher-risk components of the system. This effort is described later in this report.

HIGH-DENSITY DRAM

A Packaging Feasibility Study was initiated to identify approaches for a high-density DRAM module. The goals of the study were to provide memory density of 10 gigabits per 24 cubic inches in a cost-effective approach based on small, repairable, highly-replicated units. The approach presented is based on stacked planar modules interconnected with high-density z-axis interconnect. The module was conduction cooled to dissipate heat from the outer surfaces.

The proposed memory module is shown in **Figure 1**. This approach uses 20 stacked layers each with 124 4-Mbit DRAM chips and associated buffers. A DRAM die size assumed to be .16" x .38" results in a 4.06" x 4.24" mounting footprint. This footprint allows dice to be mounted face-up or face-down, spaced by .060" side-to-side and .050" end-to-end. This spacing allows wire bonding, flip TAB, and flip bonding die attach techniques. **Section D-D** of Figure 1 shows that substrate via pads are arranged in a grid under the die so that inter-chip spacing is minimized. Interconnect is supplied by four 100-pin connectors arranged along the short ends of each module layer. These 400 pins are assigned to power, ground, and signal categories, with some spare pins. This I/O is representative of the requirements of a 64 bit data bus, showing that this I/O will support arbitrary memory system architectures.

This straw-man module design provided a basis for a first order thermal analysis. **Figure 2** depicts the thermal model of this stacked module. For the purposes of this analysis, it is assumed that there is no heat flow perpendicular to the module layers, except at the edges through a peripheral copper gasket. The substrate material in this analysis is aluminum nitride (*AlN*). The only other material in the thermal path is the thin layer of adhesive under each die. This analysis is pessimistic, assuming the worst case where all dice consume 200mW simultaneously and the thermal path is one dimensional along the short axis of the module. The heat dissipated from the 12 dice in the center of each layer is assumed to be conducted through the four mounting holes and associated vias along the center of the long axis of the module. It is also assumed that symmetry will cause the dice to the left of center to dissipate power through the gasket on the left edge of the module and the dice to the right of center dissipate power down the right side of the module. The worst case temperature drop, from the center die on the top layer past the 6 dice between it and the edge, plus the drop to the assumed cold plate on the bottom of the module stack is depicted by the resistor equivalent of this thermal path shown in Figure 2. This first order thermal analysis shows about a 68°C temperature rise for this worst case path. Given a commercial operational die temperature limit of 100°C, the outer surface of this module would have to be maintained below 32°C (90°F).

This PFS developed a proposed approach that includes cost estimates, sources of required technology, a more complete analysis, and a mechanical mock-up.

AQUARIUS III PACKAGING STUDY

A Packaging Feasibility Study was been completed for the Aquarius III based on preliminary system specifications. Since the Aquarius III design contains devices with more than 250 pins, requiring Level I packaging, and packaging of large amounts of distributed memory at Levels II and III, the Feasibility Study provided critical information for determining implementation risks of this system design.

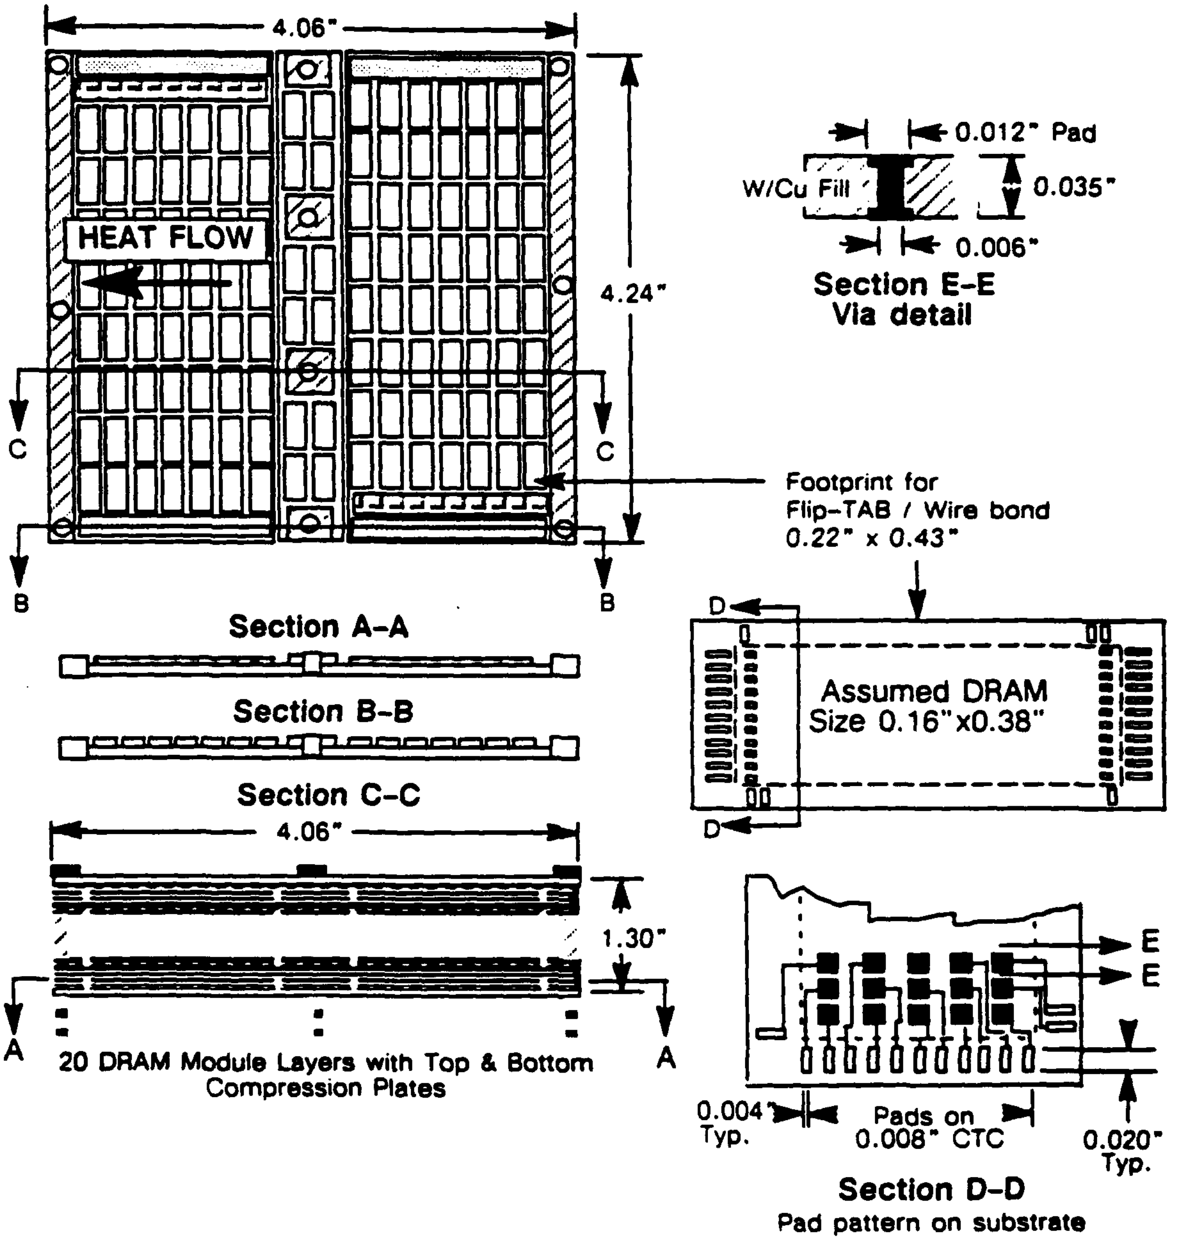

Module pin Count

72 Bits-Data I/O & Parity
124 Bits-CAS & RAS
26 Bits-Address
80 Power/Ground
98 Spares
400 Pins Total

Module Parts Count

124 - DRAMS
6 - By 10 Address Buffers
8 - By 10 BiDirectional Data Buffers
4 - 100 Pin button Connectors
(at 4 MB/Die=~10 GB)

Legend

Copper Spacer
IC Buffers
Button Connectors

Figure 1: Proposed DRAM Memory System Module

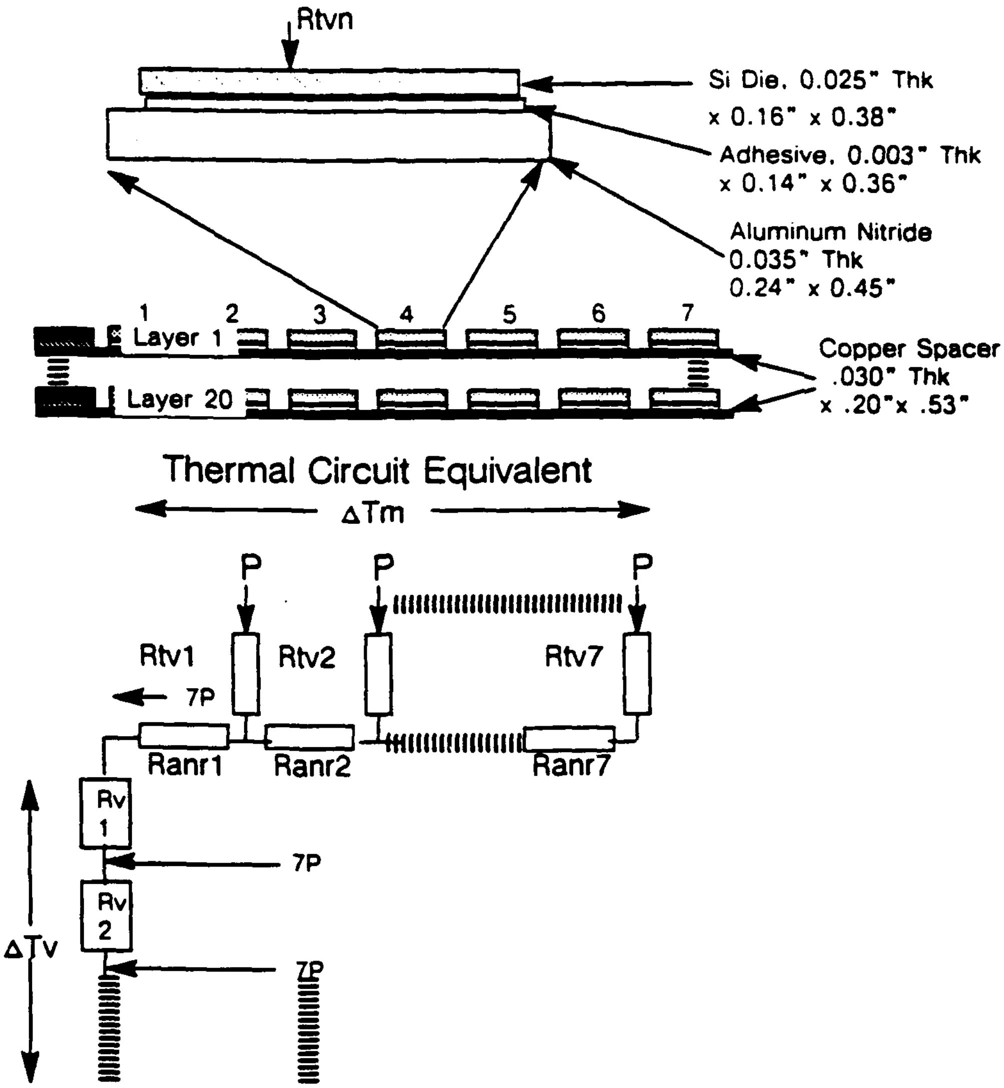

ITEM	k	Area	Length	Rt
Si Die	80	0.0608	0.025	0.06167=Rs
Adhesive	20	0.0504	0.003	0.03571=Rad
AlN(radial)	150	0.01575	0.24	1.01905=Ranr
AlN(perp)	150	0.106	0.035	0.02641=Ranp
Spacer	220	0.106	0.03	0.01543=Ros
Interface				0.35 =Rif

Rtvn=Rs+Rad=0.10
ΔTm=P*(28*Ranr+Rtv7)=3.80°C
Rv=Ranp+Ros+Rif=0.39
ΔTv=1470*P*Rvn=64.00°C
ΔTΣ=ΔTv+ΔTm=67.81°C

Figure 2: DRAM Module Thermal Analysis

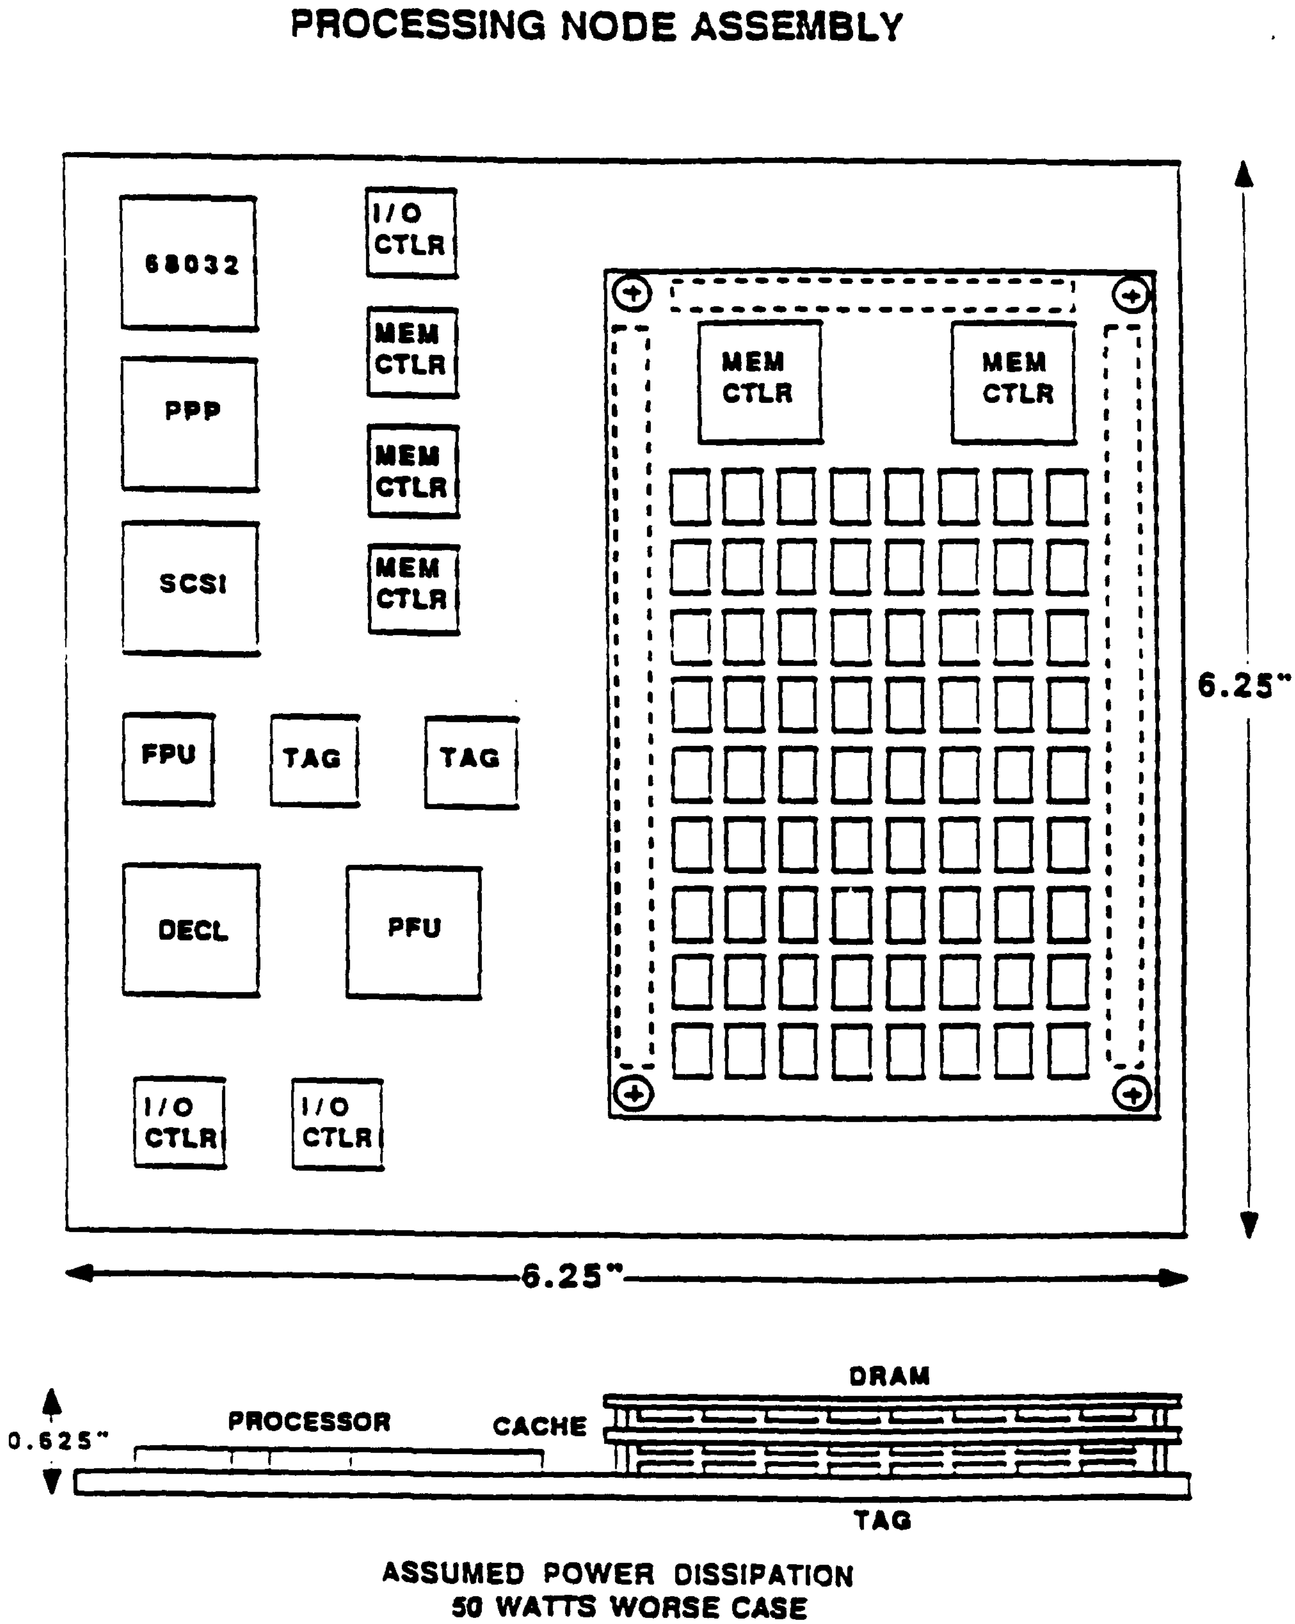

Figure 3: Aquarius III Packaging

CROSS SECTIONAL VIEW OF 1/2 OF A DUAL MOTHER BOARD ASSEMBLY

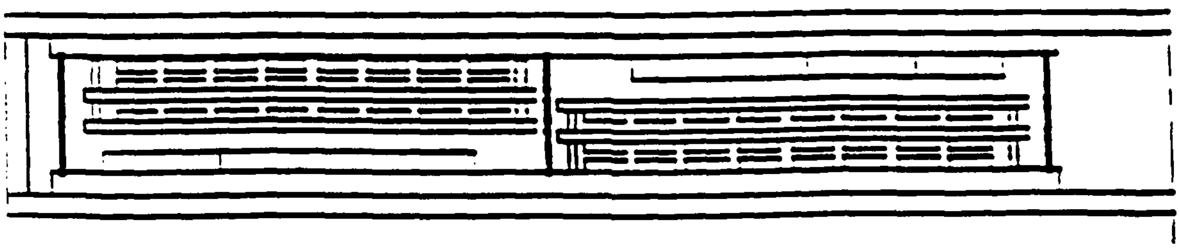

ASSUMPTIONS:
1- "Duct" diminsions - 0.25" by 3.0" by 6.0"
2- Number of ducts - 6
3- Power per duct - 32 Watts (200 Watts total)
4- Average Inlet Air Temp. (Ti) - 60°F
5- Air flow per duct - 6 Ft^3/min.

Air Characteristics:
1- Density (d) - 0.065 lbs/ft^3
2- Viscosity (µ) - 0.047 lbs/Hr-ft
3- Thermal Conductivity (k) - 0.017 BTU-ft/Hr-ft^2-°F
4- Spectic Heat (Cp) - 0.241 BTU/lb

Calculated Factors (per duct):
1- Mass flow (M) - 31.2 lbs/Hr
2- Ave. Air Velocity (V) - 9.2 x 10^4 ft/Hr
3- Heat flow (Q) - 109.2 BTU/Hr
4- Duct Equivalent Diameter (De) - 0.038ft
5- Duct x-sectional Area (As) - 5.2 x10^-3 ft^2
6- Duct wetted Area (Ad) - 0.25 ft^2

Thermal Calculations (per duct)
1- Air Temp. Rise $\Delta Ta = Q/(MCp) =$ 14.5°F
2- Reynolds Number $Re = De^*V^*d^*/\mu =$ 4835
3- Prandel Number $Pr = Cp^*\mu/k =$ 0.67
4- Duct Wall Temp. $\Delta Tm = Q/(h^*Ad)$

Where $h = 0.23^*k^*(Re)$^0.8$^*(Pr)$^0.4)/De = 7.75

Therefore $\Delta Tm =$ 56.4°F

With a 10°F Component case ΔT, $Tj =$ 10+60+14.5+56.4= 141°F

Figure 4: Aquarius III Thermal Analysis

The Feasibility Study for Aquarius III assumed that a multi-chip hybrid approach was used to implement a large machine and provided details (see **Figure 3**) on substrate sizes, die-attach methods, physical parts placement, and functional block definition. The Study also provides an analysis of system cooling requirements (see **Figure 4**) based on the proposed mechanical design, the number of anticipated nodes, and the power consumption of each node.

MIT ALEWIFE

The Phase I MIT Alewife system is a 64-processor multicomputer based on *Sparcle*, a prototype processor derived from LSI Logic's SPARC implementation. Sparcle clocks at 33 MHz, resulting in a peak throughput of 2 GIPS for a 64-node machine.

The Sparcle processor uses a custom memory controller to hold cache tags and implement cache coherence protocols by synthesizing messages to other nodes. A control word associated with each memory reference allows various synchronization or communication data types to be synthesized by the controller. The controller signals to a remote memory module when a processor context switch has been caused by a synchronization fault or a cache miss.

Besides the processor and memory controller, each node has 64K bytes of direct-mapped cache and 8M bytes of main memory. The memory on each node is partitioned into a 4MByte globally-shared portion, and a 4MByte local memory part, a portion of which is used for the coherence directory. Thus a 64-node machine has 0.5 gigabytes of memory. A numerical co-processor and a Frontier series Mesh Routing Chip (FMRC) from Caltech comprise the rest of the node. Free ports on peripheral nodes of the network are used for I/O, monitor, and host connections. The prototype Alewife system will attach to a host SUN by interfacing a network switch to the VME bus.

PROTOTYPING AND PACKAGING - PROPOSED APPROACH

PROCESSOR CARD

The entire circuit of each Alewife processor node - CPU, FPU, MMU, FMRC router, and memory - was contained on a commercial form-factor card. This approach was taken as none of the devices are available in unpackaged form. The combination of device packages (PGA, PLCC, SOJ, and TSOP) requires that a mixture of assembly techniques be employed to produce the first group of prototypes.

Each processor card was fitted with a standard commercial "pin-and-socket" connector to provide interconnection to the backplane. The connector will carry the mesh signals required for the node, plus pins for system-level signals such as clock and reset. Pins were allocated as needed to card power connections.

BACKPLANE

The processor cards plug into custom backplanes designed to fit standard cabinetry. This allows multiply-sourced, off-the-shelf mechanical components to be used wherever possible to reduce implementation costs and procurement time.

The backplane was designed with 20 card positions. This allows the backplane to house a single 16-processor row from the 2D routing mesh and provides additional slots at each end for host interfaces or controller cards for peripherals. Systems smaller than 16 columns in width can be constructed by using clusters of adjacent slots in the backplane.

COMMUNICATION MESH ROUTING

The FMRC routers rely on short wiring lengths to achieve high signalling rates. In this proposed approach, the inter-FMRC wire length can be maintained at 4 inches or less for routers within a backplane, and 6 inches or less for routers signalling between backplanes. The anticipated round-trip time for FMRC transfer control signals traversing these wires is approximately 2.5 nanoseconds, implying that the impact of this approach on inter-node communications is minimal for a 128-node system (16 columns by 8 rows).

SYSTEM EXPANSION

Connectors were provided at the top and bottom edges of each row backplane, allowing a vertical array of backplanes to be tied together via ribbon cables. Additional connectors at the left and right edges of each backplane allows adjacent racks of backplanes to be joined. This approach allows Alewife systems of arbitrary size to be built.

SYSTEM CABINET AND COOLING

The Alewife system was housed in a commercial cabinet. System cooling issues are largely eliminated in pre-engineered commercial housings which guarantee that an airflow of 400 LFM is maintained.

SYSTEM POWER

Worst-case power requirements for an individual processor card are estimated to be approximately 5A @ 5VDC. Overall worst-case power requirements are thus approximately 320A @ 5VDC, or 1600 Watts. Specifications indicate that supplies in this power range are equipped with their own cooling fans, further simplifying the issues of cooling.

EXPERIMENTAL PACKAGING APPROACH

INTRODUCTION

The 2D routing mesh of the Alewife system creates an opportunity for investigating unconventional approaches to systems packaging. This experimental package eliminates the system backplane by interconnecting the processor nodes via z-axis "stacking" connectors.

UNITS OF REPLICATION AND SCALING

In the experimental Alewife, the level of system modularity is the processor card. A 2D Alewife mesh of arbitrary size can be implemented by producing the appropriate number of processor cards. Standard 2D meshes require an array of one or more backplanes, which in turn require cables to provide communication across the backplane seams.

COOLING

Cooling the experimental Alewife is simpler than a backplane-based system. The cooling airflow path is not shared by multiple cards as in a backplane system, eliminating potential hot spots created by pre-heating the cooling air as it passes over several cards. It also reduces drag in the channels so that cooling fans can be smaller, generating less noise.

PROCESSOR BOARDS

In the experimental package, boards (**Figure 5**) are notched to allow them to "key" on the rails of the board compression frame. The notches are offset to correctly orient the boards in the frame. Plastic blocks are placed above and below the connectors to maintain the airflow channels. Each block also contains an alignment pin that "keys" the cards together. Since the connector contact alignment is not sensitive to board misalignments of up to 10 mils, the keying arrangement can consist of rounded or tapered pins protruding into the next board.

BOARD INTERCONNECT

The experimental package for the ALEWIFE requires that the processors be interconnected by means of z-axis plunger contact connectors. These connectors would be attached to the boards with flat-head screws, allowing simple connector replacement.

The connector envisioned for this experiment, fabricated by Augat. Inc., uses pad area interconnect (PAI) contacts and an inert material for the contact shell. A cross-section of a PAI contact is shown in **Figure 6**. Accommodating the 2-D FMRC data paths and adequate power current routing through the board array would require a modified version of an existing 210-contact connector. In this device, the contact rows are 50 mils apart and the contacts are placed on 100-mil centers within the rows.

As shown in **Figure 7**, the dimensions of this z-axis connector are approximately 0.450"(W) x 3.5"(L) x 0.437"(H). These connectors would be attached to the processor cards with flat-head screws threaded into the connector body.

SYSTEM CABINET

The ALEWIFE cabinet would be custom-fabricated to contain the components of the system: board array compression frame, power supply, and cooling fans.

BOARD COMPRESSION FRAME

The board compression frame proposed for the experimental package (**Figure 8**) would provide the force to fully compress all connector contacts.

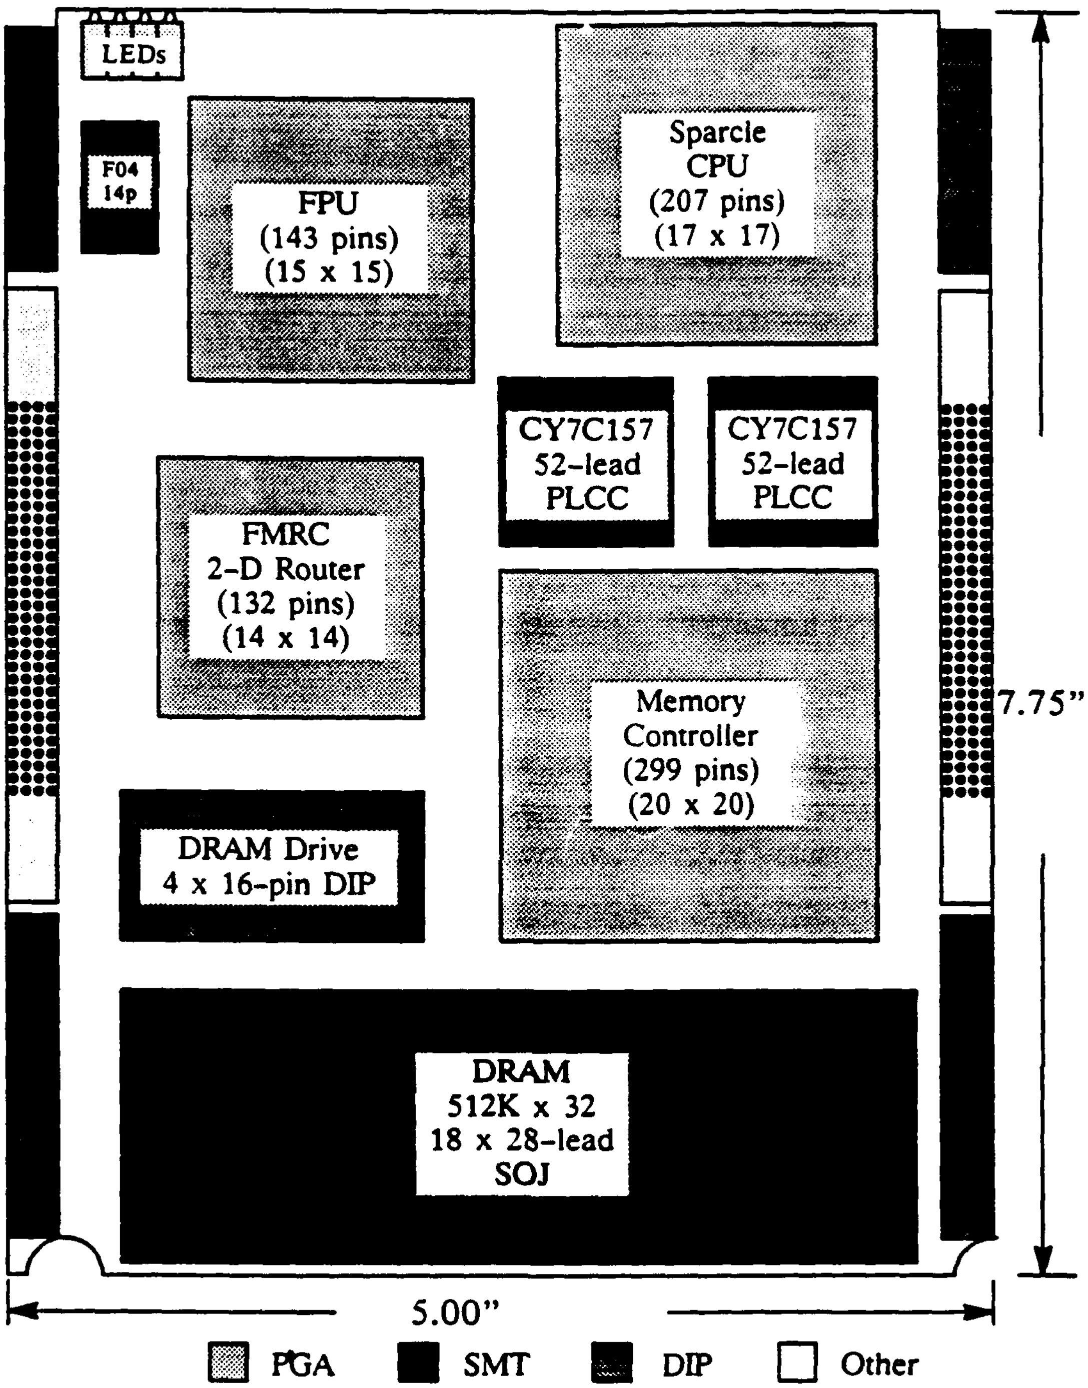

Figure 5. ALEWIFE Processor Card (experimental package)

The compression frame will equalize force on the primary compression areas by means of springs. The springs also minimize the effects of accumulated tolerances in the z-axis.

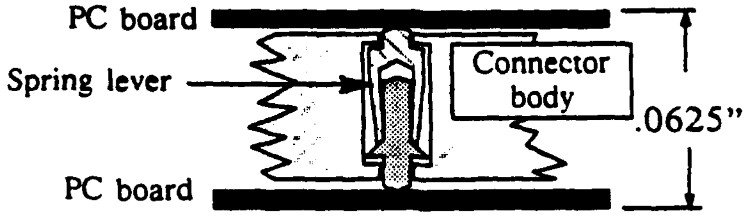

Figure 6. Augat PAI Contact

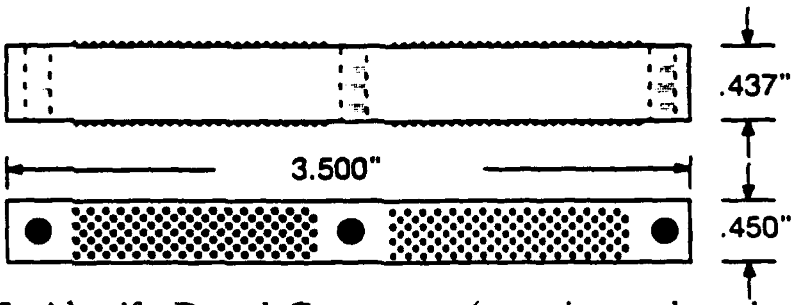

Figure 7. Alewife Board Connector (experimental package)

SYSTEM COOLING

COOLING FANS

In the experimental Alewife package, system cooling consists of seven 7" Rotron fans blowing in the upward direction. After allowing for back-pressure, an estimated air velocity of 400 feet/second is expected in the inter-board channels. In an office environment of 25°C, this airflow should be more than sufficient to maintain the VLSI junction temperatures at or below 65°C. This operating point is low enough to insure reliable operation.

STATUS

A mechanical mock-up of the experimental Alewife system package has been completed.

Collaborative Development Efforts

INTRODUCTION

Collaborative Development Efforts (CDEs) are the most involved and advanced level of service provided by APT. A CDE involves the APT engineering staff in the design and fabrication of a critical part of a new architecture to investigate new uses of advanced packaging technologies for improved performance. Typically conducted in parallel with the clients' primary development efforts, these experiments provide the collaborating partner with low-risk, directly relevant packaging demonstrations that can be folded into the product cycle. APT CDEs are reviewed below.

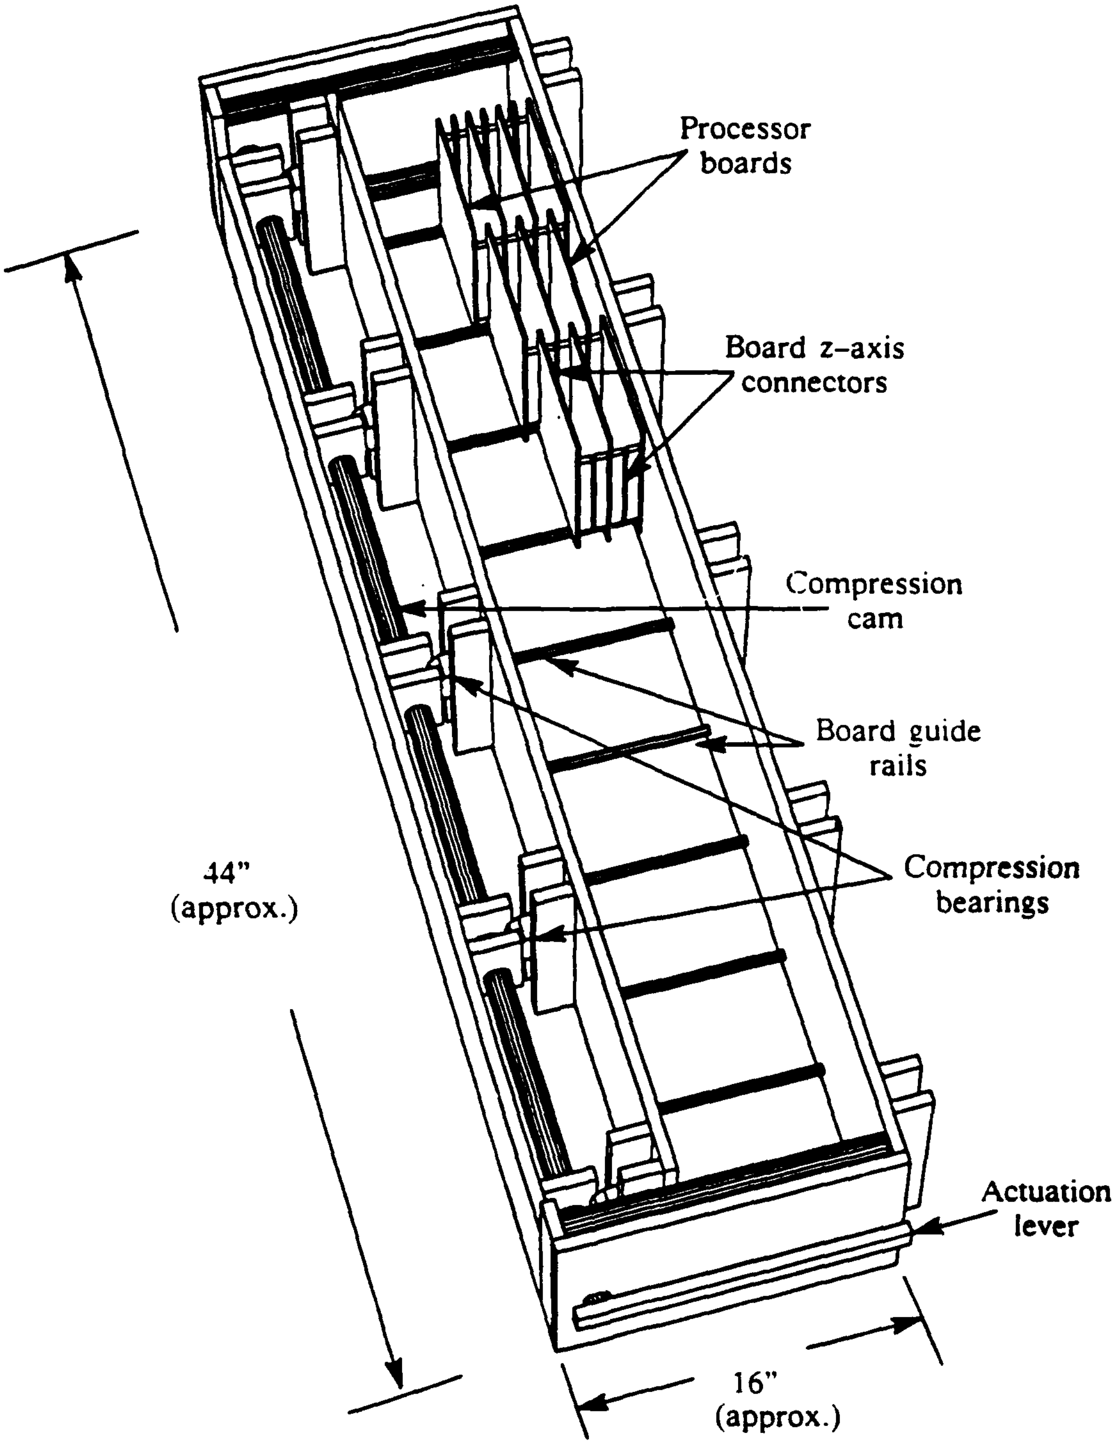

Figure 8. Alewife board compression frame (experimental package)

BBN SWITCHING TECHNOLOGY DEMONSTRATION

Under contract to DARPA, BBN developed the Monarch, a parallel architecture scalable to thousands of processors connected to a large, shared-memory system. A critical element of the Monarch architecture was the interconnection network through which processors access the memory system. In the Monarch Medium Scale Prototype designed by BBN, the interconnection network was packaged with processors and memory system components on printed circuit boards, joined by a crosswise interconnection structure.

Using an additive copper-polyimide substrate technology, APT packaged a section of the Monarch interconnection network as a multi-chip module, referred to here as the Switch-Concentrator Module (SCM). A Monarch interconnection network of arbitrary size can be built by connecting multiple instances of this module.

The SCM is an excellent choice for a packaging technology demonstration because it requires high-density interconnect, significant power dissipation, controlled-impedance transmission lines, and high edge speeds. In addition, this module is a well-defined component that can be easily extracted from a larger system. Experimental results from this high-risk technology may be directly compared to results from an implementation that uses more conservative packaging technology.

High-density wiring and direct bonding of chips to this polyimide substrate enabled packaging of each module in a fraction of the volume required by conventional printed circuit board technology. Low-dielectric insulators and controlled-impedance properties were designed to support maximum signal rates of the silicon design. Thermal design supported the 100-plus watts expected from this module.

OBJECTIVE

This demonstration of custom VLSI signaling technology and ISI-designed substrate was supported by an ECL-based test setup and test software running on a SUN workstation. Specifically, the objectives were:

- Demonstrate viability of packaging technology for high-performance systems.
- Verify high-speed VLSI signaling technologies developed for Monarch.
- Verify substrate transmission line quality using high-speed ECL devices.
- Characterize transmission lines using time domain reflectometry (TDR).

WATER-COOLED HEAT EXCHANGER

A water-cooled heat exchanger was designed and fabricated. This exchanger was intended to avoid overheating of the ICs and substrate if the heat-transfer structures of the SCM proved to be inadequate. Water cooling was chosen to maximize heat transfer and keep the test environment clean and stable (e.g., vibration- and dust-free). The exchanger was driven by an aquarium pump.

DEVICE DE-PACKAGING

The number of unpackaged devices available for pre-assembly screening was quite small (under 40); the expected yield from testing was at best six devices. BBN had 47 known working packaged devices; it was decided to "de-package" these chips and use them as the initial pool for device screening.

The de-packaging operation was performed in two steps. First, the bonding wires were cut off the devices, leaving the original wedge bonds in place. Second, the packages were heated to approximately 120°C, causing the die-attach epoxy to release and allowing the devices to be scooped out of the package. 45 of the 47 devices mechanically survived both operations.

DEVICE SCREENING

ISI had an IC probe card fabricated to connect essential signals, supplied by the test apparatus, to an unpackaged device on a semi-automatic probe station. Screening was conducted by ISI at 10 MBaud on eight of the untested dice. None of the devices fully passed the suite of acceptance tests. The same tests were performed on twenty-four of the de-packaged devices, producing results ranging from total failure to complete functionality.

Two of the screened devices passed all tests and were used as the basis for the signaling demonstration. The remaining screened devices were sorted according to degree of functionality. Anticipating that a good device might not pass the screening process due to poor-quality signals delivered by the probe card, all partially functioning devices were reserved as backup in the event of failure of the primary devices.

SUBSTRATE PRE-ASSEMBLY AND TEST

The ECL clock distribution devices and discrete components were assembled onto the substrate in one operation. The ECL devices were attached with non-conductive epoxy, while the discrete components were attached using EPO_TEK E20 conductive epoxy (80% silver loaded). The epoxies were then cured in an oven for one hour at 70°C.

The substrate wiring was modified to bypass the diagnostic "daisy-chain" of a fully-populated SCM. This was done by wiring across the low-speed interconnects with 30-gauge wire.

The substrate was attached to the test setup and the clock distribution network was debugged. Several forms of failure were found in the substrate interconnect during this process. Repairs were effected in the most expeditious manner using 50 ohm wire-wrap coaxial cable and, where needed, bonding wires strung like "telephone wires."

The debugged substrate was returned to the assembly house for mounting and bonding of the switch devices. The dice were attached to the substrate with conductive epoxy, which upon curing formed the thermal and device-substrate ground connections.

SIGNALING TECHNOLOGY DEMONSTRATION

The assembled substrate implemented a 2-chip circuit demonstrating the essential features of Monarch inter-chip communications. The test system clock was generated by a 300 Mhz pulse generator, which allowed multiple operating frequencies to be investigated. A small, wire-wrapped 100K ECL circuit produced appropriate system clocks, data frame transmission synchronization signals, and clocks for the low-speed serial diagnostic bus.

The waveforms captured in this demonstration generally support the ISI belief that the BBN 1.6 micron CMOS devices are designed with sufficient signaling headroom to allow inter-chip communication at peak data rates of 300 MBaud.

MECHANICAL STUDY

ISI commissioned an independent study of the failures observed in the substrate to determine the nature of failures seen in the switching demonstration. The mechanical study also provided a physical basis for results observed during the TDR analysis. Results from substrate sectioning are presented here.

The substrate failures were all attributed to separation of plating interfaces under z-axis tension during thermal expansion of the polyimide dielectric. This failure mode can be traced to two sources: inadequate specification of design rules, and inadequate fabrication process control. The original design rules made no restriction on the ratio of surface pad area to column area at layers deeper in the structure.

The failure mechanism was discussed with the manufacturer; fabrication process modifications were made. Further refinements in processing are mandated by the poor adherence to nominal feature geometric tolerances.

THERMAL MANAGEMENT

Removing heat from high-density systems remains one of the most formidable challenges in system packaging. Therefore, the thermal conduction capability of multi-chip module substrates and single-chip packages is of considerable interest to system designers.

The SCM substrate was designed to conduct a moderate amount of heat (4 watts per device) away from the switch and concentrator devices with only a mild rise in device operating temperature. The mechanism used to perform this conduction is called a "thermal column," and is shown in **Figure 9**. Sixteen thermal columns are used for each device.

Using the substrate described earlier, a thermal conduction study was conducted by simultaneously sampling the surface temperature of one of the VLSI devices and the substrate cooling surface immediately below the devices while the substrate was thermally isolated from the probe station chuck. To avoid an upward frequency drift of the HP generator observed at high operating frequencies, and to maintain VLSI operating stability, measurements were taken at transfer rates of 170 MBaud. With the VLSI devices operating at 170

MBaud, the supply current was measured to be 1.5A, indicating that the devices were dissipating approximately 3.75 Watts each. The setup for thermal measurements is shown in **Figure 10**; the thermal measurement equipment consists of calibrated Chromel Constantan 0.020" diameter thermocouples and Keithley Model 197 digital microvoltmeters.

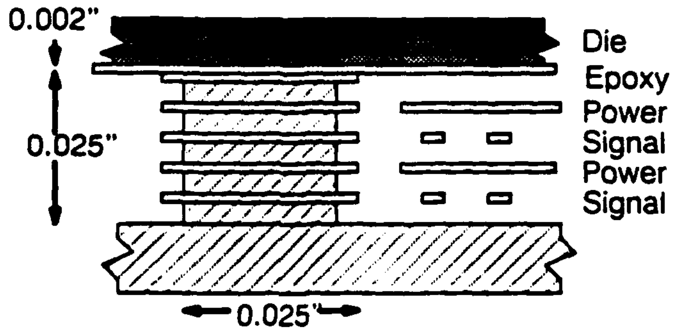

Figure 9: SCM thermal column

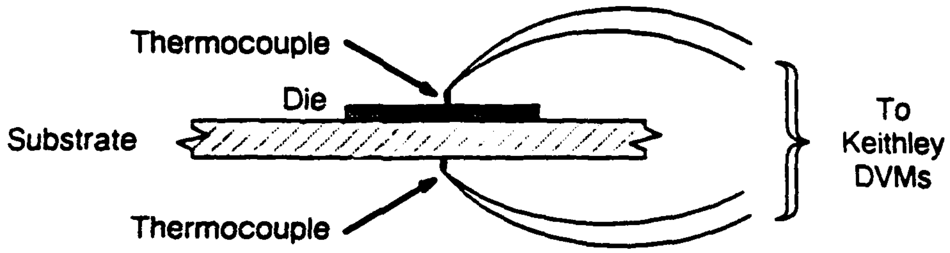

Figure 10: Thermal measurement setup

The voltage readings from the thermocouples, shown in **Figure 11**, were interpreted using tables provided by the manufacturer. The abrupt initial temperature rise of the device was due to the thermal resistance in the device-attach epoxy joints and the substrate thermal columns. Once the epoxy and thermal columns began to conduct, the entire substrate conduction slab began to warm, roughly tracking the device temperature. At a remote point in time, not shown here, equilibrium was established, and convection cooling of the slab maintained a constant device temperature of approximately 46°C.

A preliminary thermal analysis, performed in 1988, predicted that worst-case power consumption (4 Watts per VLSI device or 64 W/in^2), would cause device junction temperatures to rise a maximum of 22°C above the cooling surface (back) of the substrate, or 5.5 °C/W. At 3.75 Watts each, the actual rise is approximately 10°C, which equates to thermal resistance of 2.67°C/Watt, including die-attach thermal junctions. The substrate was operated in this manner for all testing; it was decided that there was no need to complicate the test setup with an unnecessary water-cooling plate.

TRANSMISSION LINE EVALUATION: ECL TRANSMISSION

Because the 1.6 micron BBN devices displayed very fast signaling potential, the possibility of examining signal quality at frequencies of 300 MHz and higher warranted separate inves-

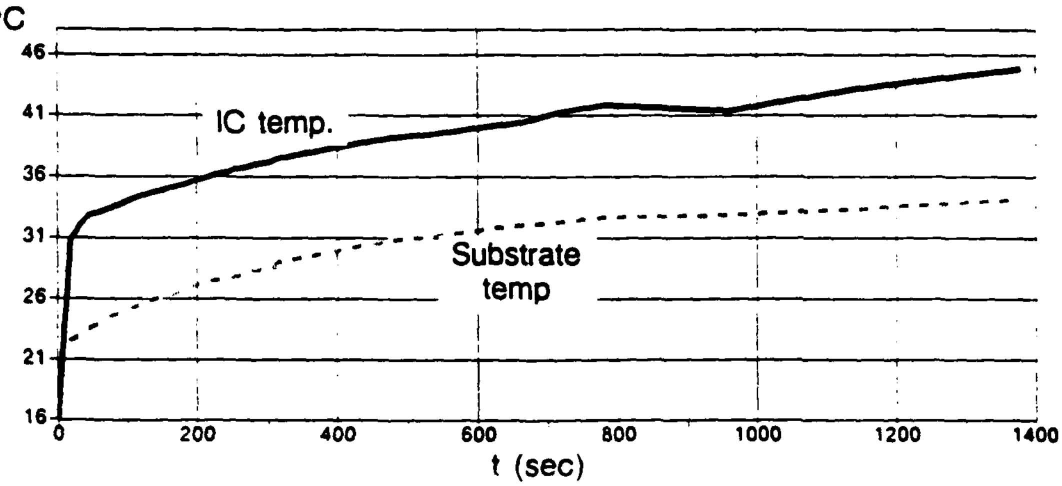

Figure 11: IC and substrate temperature

tigation. Motorola MC100E111, very fast 1:9 clock buffers, were procured to drive high-speed waveforms onto the substrate clock distribution lines.

The unpackaged MC100E111 ECLinPS ("eclipse") devices were attached to a second substrate using non-conductive epoxy and conventional wire-bonding techniques. The clock generator used for the VLSI signaling demonstration was connected directly to the substrate.

The results of the TDR study show that the line quality produced by the additive polyimide-copper technology is highly suited for use in high-speed systems. The inter-line coupling characteristics are excellent. Process controls must be improved, however, for substrates to be used in 50Ω systems where back-reflections caused by incorrect line impedances (and thus improper termination) cannot be tolerated.

POLYIMIDE TECHNOLOGY EVALUATION COUPON

The fabricator was consulted on the feasibility of adding a small test coupon to the SCM substrate design. There was room to fit a 0.5"x6" test area next to the SCM and still meet clearances for a repeated pattern on a standard panel.

The intent of the test coupon was to evaluate the quality of the transmission lines produced by the standard Microtec process. To this end, the following structures were included in the design for evaluation with a time-domain reflectometer (TDR). The process evaluation structures were all 50-Ohm coaxial connections to transmission lines terminated by surface-mount chip resistors.

A short stripline transmission line. Unperforated ground planes are atypical (and very difficult to fabricate) in the Microtec process.

A transmission line between perforated ground layers, typical of Microtec substrates.

A transmission line containing a single 45-degree bend.

A transmission line containing a 90-degree bend.

A transmission line containing two standard vias.

Two parallel lines at the closest allowable spacing, allowing measurement of inter-line coupling.

Three lines characteristic of geometries found on the SCM layout.

A very long (2 foot) transmission line with segmented corners, intended to demonstrate the loss characteristics of standard lines.

SUMMARY

The APT effort to develop an alternative packaging approach for the BBN Monarch was completed. Results were demonstrated in the following areas:

MULTI-CHIP SUBSTRATE TECHNOLOGY

The additive copper-polyimide process in principle provides a viable alternative fabrication technology for multi-chip modules needed for high-performance systems. Results show that the thermal management properties are excellent. Signal transmission and inter-line coupling properties are very good.

The mechanical study of substrate failures indicates that drawbacks to the current state of the art are primarily due to inadequately specified design rules and lack of controls in the multiple stages of the fabrication process. Additional work on fabrication process dimensional control is warranted to eliminate the deviation from desired line impedance.

SWITCHING TECHNOLOGY

The results of the signaling demonstration show strong evidence that BBN has made a significant advance in the state of the art in inter- and intra-chip signaling rates.

<u>ENCORE COMPUTER - GENESIS</u>

INTRODUCTION

Under DARPA sponsorship, Encore Computer designed the Multimax II, an enhanced version of the Multimax parallel processor. Multimax II is based on the Motorola 88000 RISC microprocessor. ISI and Encore jointly produced Genesis, a demonstration version of the Multimax II system, using advanced packaging technology developed by APT.

The Genesis project employed several packaging technologies: TAB packages, button/plunger interconnect, and fine-line PCB. These technologies were used to develop a High-Density Systems Module (HDSM) processor and cache sections of the Genesis demonstration.

ISI's HDSM format is based on fine-pitch, stackable connectors to interconnect layers of multichip modules, which can be configured to meet different systems requirements. Specifically, APT has developed a packaging technique that can be used to implement both generic and application-specific MCMs that can be stacked and mounted on a printed circuit board, or built into frames that can be stacked in three dimensions. The HDSM-based Genesis

module, comprised of two processor module layers and one memory module layer, is shown in **Figure 12.**

GENESIS PROCESSOR

The Genesis processor is a quad-88000 CPU designed to fit onto two modules fabricated in the ISI High-Density System Module (HDSM) format. Each module contains two 88100 CPUs, four 88200 Cache Memory Management Units (CMMUs), and 22 discrete components. The 88000 devices are available from Motorola in a variety of forms.

To maximize the packaging density of the Genesis module, negotiations were conducted to procure the 88000 devices in die form. Motorola agreed to supply tested, burned-in 88000 components mounted in JEDEC standard 188-lead TAB frames.

The selection of an appropriate substrate technology for the processor module was driven by the physical dimensions of the TAB package and the topology of the dual-88000 processor

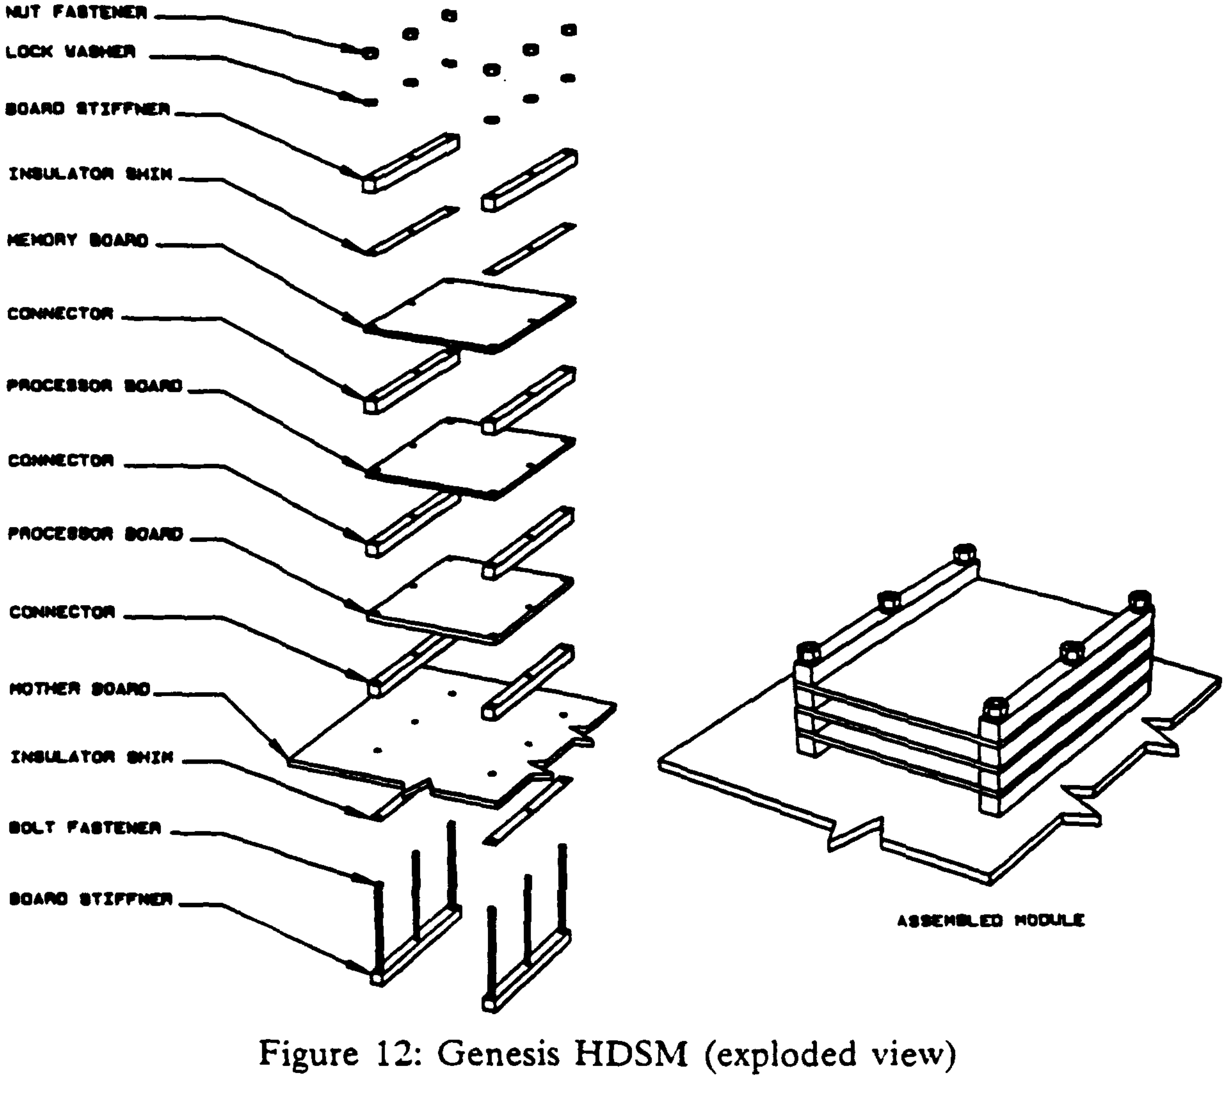

Figure 12: Genesis HDSM (exploded view)

circuit. The 188-lead package required a substrate footprint approximately 0.900" square with 9-mil-wide pads on 15-mil centers, spaced evenly around the perimeter. A 0.450" square copper area in the center of the footprint pattern is required for die mounting and heat transfer. The processor module top layer layout is shown in **Figure 13**.

Examination of the processor circuit topology indicated that 4 layers of fine-line printed circuits would be required to interconnect the 88000 components within the HDSM format. More exotic - and expensive - MCM substrate technologies would not be required. Accordingly, the processor substrate was designed using 4-mil design rules.

A design requirement for 62 Ohms characteristic impedance mandated that 4 power/ground planar layers be included in the design. The planar layers are alternated with the signal routing layers, forming striplines on the inner layers and microstrips (embedded in solder resist) on the outer layers.

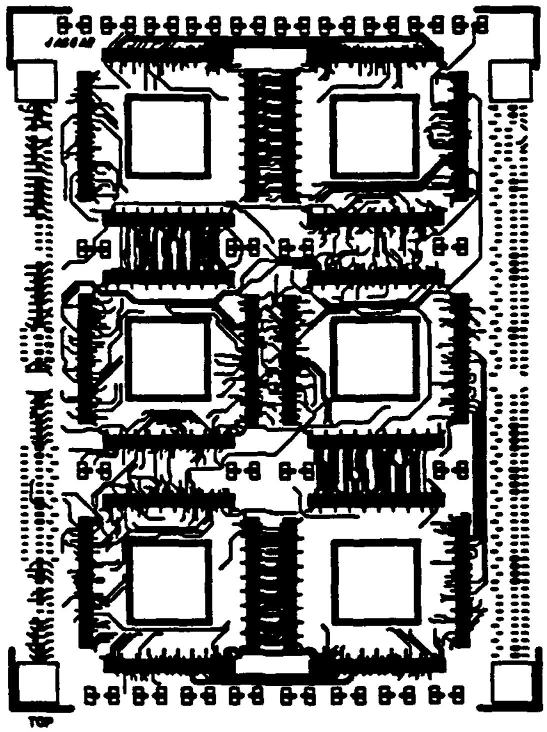

Figure 13: Genesis HDSM Processor Top Routing (1X scale)

Alternating substrate layers, ordered as shown in **Figure 14**, act to reduce interconnect crosstalk by eliminating inter-layer coupling. The central Vdd and Gnd planes, separated by a 5-mil dielectric layer, develop a 10nF decoupling capacitor. The planes are fabricated from 1.4-mil-thick copper, producing an extremely low internal resistance and self-inductance capacitor that is very effective in reducing high-frequency switching noise.

Assembly of the processor module TAB components was completed with the cooperation of Motorola. Five of of the completed modules were tested with only one device failure found.

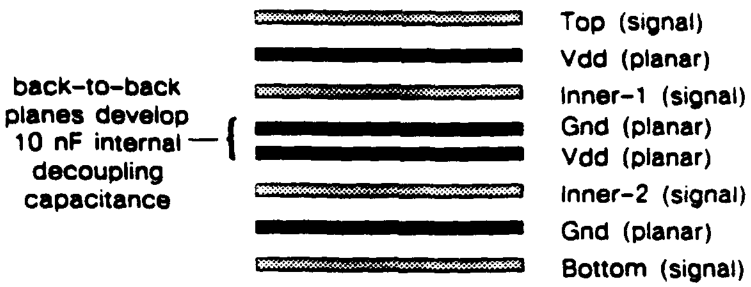

Figure 14: Genesis HDSM Processor Substrate Layers

Testing was conducted at ISI using the ISI/MCC ES-Kit adaptor board described later in this report. Test software consisted of modified ES-Kit programs including power-on diagnostics, a simple monitor/loader, and a demonstration program written at ISI.

PROCESSOR DIE ATTACH

A number of liquid die-attach media are available. However, liquid materials may be difficult to control in both handling and application, particularly in low-volume prototyping situations. Since the 88000 devices are fabricated with p-well and n-well processes, these devices require a guaranteed electrically insulating die-attach. To avoid pushing the dice through a liquid epoxy, it was decided to use a "B-stage" or partially-cured epoxy film, as the die attach medium. A number of vendors were contacted, with the result that A.I. Technology TK 7758 3-mil-thick epoxy film adhesive was selected.

The TK 7758 material is aggressively tacky when initially applied, creating some device handling problems during assembly. When cured, however, it forms a mechanical buffer for the stresses arising in the different thermal expansions of the processor substrate and the processor devices. In addition, it guarantees that there will not be any voids in the die-attach, forming a reliable electrical insulator for die substrates. This die-attach method has not failed on any of the 42 processor devices that were mounted, and has been adopted by Motorola for internal use after observing ISI assemblies.

The TK 7758 material is loaded with aluminum nitride (AlN) to provide excellent thermal conduction characteristics. The expected thermal resistance of this module is approximately 2°C/W.

GENESIS MEMORY

Along with the two dual-88000 processor modules, the Genesis HDSM stack contains the node cache memory. This memory module contains 42 high-speed static RAM (SRAM) devices to implement 1MByte of data cache, plus data parity, tag, and state memories. Since dense, high-speed memories are of general interest to systems developers, considerable effort was devoted to this module.

Tremendous leverage is gained from densely packed memory subsystems as they are usually replicated many more times in a system than processing elements. The intent of the Genesis

memory was to investigate packing densities approaching wafer-scale. Three separate designs were completed for dense memory systems. The final implementation was the result of a design cycle, shortened by schedule constraints, to rapidly prototype a surface-mount memory module.

SURFACE-MOUNT MEMORY

A number of implementation technologies were investigated for the memory module. From a baseline design for the data cache and cache tag memories, an exploration of conventional and advanced packaging technologies was pursued. For the final implementation a rapid prototyping effort was initiated to deliver a memory module to Encore.

A high-density, double-sided surface-mount version of the memory module was developed to conform to the HDSM format. Drawing on the baseline design, the module schematics were created in one day and the substrate layout was completed in less than a week. Components were procured from FAST while board fabrication was in progress.

This design uses conventional 2-sided surface mount technology to house the required components. This is an extremely cost-effective package and demonstrates the flexibility of the HDSM format in accepting a mix of packaging technologies while meeting physical and electrical design requirements.

ALTERNATE MEMORY TECHNOLOGIES

Several additional high-density memory module implementations were investigated. Each concept was taken through the substrate layout process, with the designs targeted for the fabrication process of Unistructure Inc. of Irvine Calif. Unistructure was selected as the packaging fabricator for cost reasons; the anticipated impact of substrates with excellent packaging densities and thermal properties at a fraction of the cost of other MCM substrate vendors was judged to be worth the evaluation effort.

As each design was developed it was reviewed at Unistructure to insure compliance with the fabrication process. The reviews twice uncovered problems with miscommunication of the Unistructure fabrication design rules; hence each review motivated extensive redesign to submit a design.

MULTI-CHIP FLAT LEADLESS CARRIERS

The most aggressive alternative approach for the Genesis memory module involved the development of Multi-chip Flat Leadless Carriers (MFLCs). The MFLC concept produces multiple-die modules with pad-area-grid contacts that can be mounted on a substrate to develop a packing density within 5% of wafer-scale. Unistructure proposed to fabricate these modules using individual devices. Several unsuccessful attempts were made to develop these modules for Genesis.

Unistructure used a variation on their additive MCM substrate fabrication process to develop the MFLCs. Instead of beginning with a bare aluminum plate, unpackaged memory de-

vices were arranged face-up in closely spaced groups on a flat plate. The dice then had copper bumps grown on their bonding pads to provide attachment points for the FLC interconnect. The die were then passivated and a multilayer copper/polyimide interconnect structure was fabricated directly on them. As shown in **Figure 15**, each FLC provides the electrical interconnect and mechanical housing for groups of 4–8 dice. Of particular interest is the concept of "pad relocation," or changing the interconnect pitch from 4-mil centers typical of dice to 15- or 20-mil pitch typical of device footprints found on substrates.

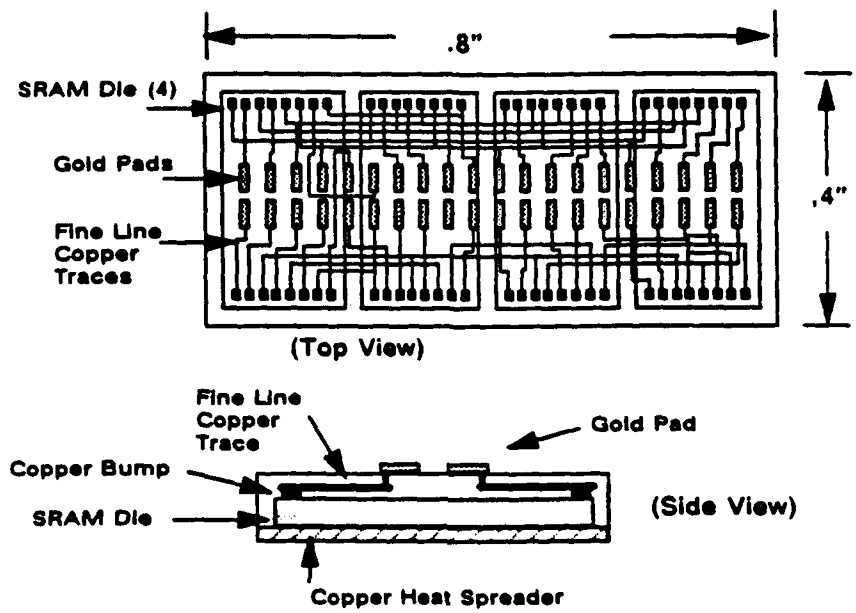

Figure 15: Multi-chip Flat Leadless Carrier

A fundamental problem was encountered with the planarity of the polyimide in the spaces between die. During curing, a "dimple" would form in the polyimide in the inter-chip areas. These surface irregularities could not be planarized and this prevented the essential upper layer metalization steps.

During the evaluation phase of this multi-chip FLC approach, ISI commissioned an independent test and analysis lab to analyze cross-sections of the processed FLC structure. Of particular interest was residual organics in the processed parts, unusual metallurgy in the metallic boundaries, mechanical alignment of the process steps, and general integrity of processed structure.

The analysis was performed by an outside laboratory and a complete report presented to ISI. The only anomalies encountered in the analysis were process residuals – a thin layer of metal on the die surface and some organic solvents that had not escaped during polyimide curing. The registration, metallurgy, and integrity of the structure was found to be adequate. From the analysis, ISI concluded that the basic process technology was sound and that it was

likely that the only problem blocking successful MFLC production was the planarization issue.

SINGLE-CHIP FLCS

After the mechanical problems with the MFLC fabrication process were discovered, an attempt was made to develop a process for individual chips to create SFLCs, or Single-chip FLCs. The SFLC packaging approach approximates the high packing density of the MFLC - devices may be spaced within 0.010" of each other by means of a metal web, or shim.

Since the pad relocation interconnect built on each of the individual die does not protrude into the inter-chip "alleys" of wafers, the SFLC is very well suited to wafer-lot processing and allows standard sawing techniques to be used for dicing. It is well established that tooling and fabrication costs may be shared by processing wafer-lots. Cost estimates given by Unistructure indicate that, in production, SFLCs would cost $1-$5 each, far less than standard ceramic packages.

Problems again surfaced in the SFLC fabrication process. Once the pad-relocation process was completed, the devices were tested. When none of the parts passed the electrical screening, physical examination of the completed parts were performed. It was found that an early process step attacked any aluminum metal of the bonding pads that was not covered by the copper deposition. This corrosion effectively disconnected the devices from the outside world.

When the SFLC process was shown to have chemical as well as die-alignment problems, further FLC development was deferred until after the Genesis memory module could be delivered. In the meantime, a complete SFLC-based memory module design had already been completed, reviewed at Unistructure, and documented. The layout produced for this design is shown in **Figure 16**.

CONDUCTIVE ELASTOMER INTERCONNECT

The MFLC and SFLC modules were to connect to their substrate by means of anisotropically conductive elastomers. Elastomers are rubber or plastic sheets that are filled with metal filaments placed at regular intervals. The filaments are oriented such that electricity is conducted along one axis only, usually the thinnest dimension of the elastomer.

Since this interconnection is made by compressing the elastomer between the FLC and substrate contacts, the resulting module would be easily repaired by disassembling the compression mechanism. Furthermore, since one of the compression plates pushes directly against the back of the devices, forced-air cooling of such a module is trivial.

TAB DEVELOPMENT

Among the most robust permanent, dense die-attach processes is face-down TAB, or "flip-TAB." Again, the low tooling cost, relatively fast turnaround, and short-run capabilities

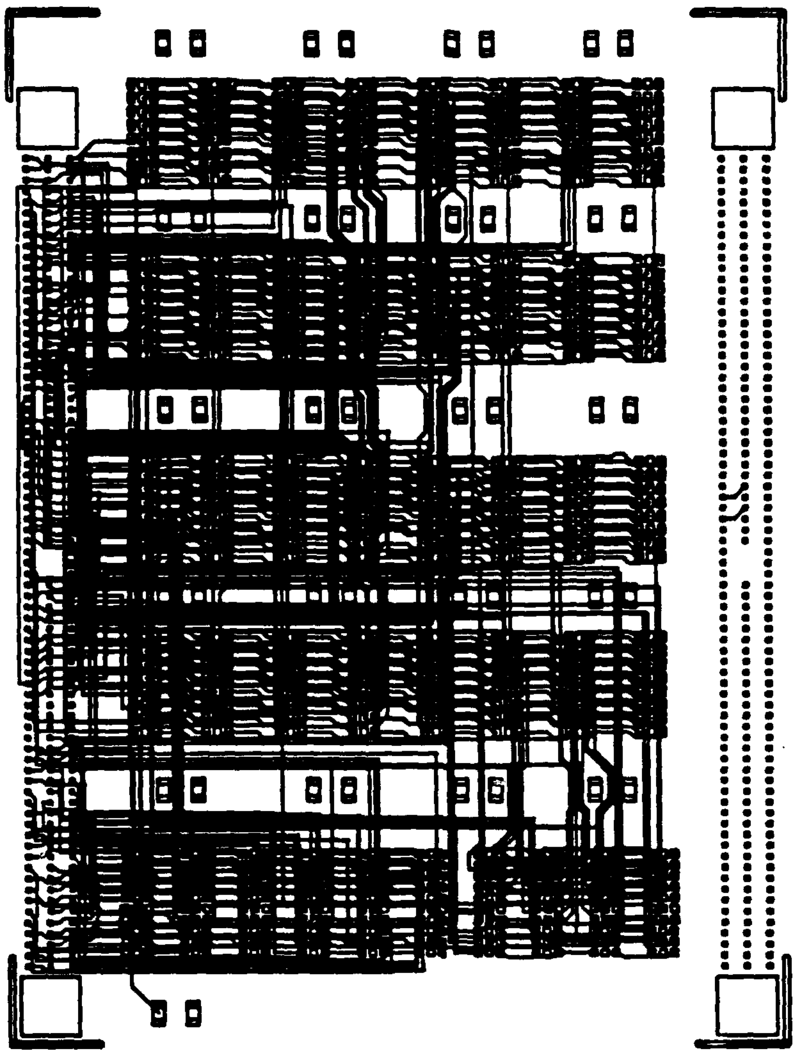

Figure 16: Encore Genesis Memory – SFLC Version (1.5x scale)

offered by Unistructure were attractive. The Unistructure "maskless" tooling approach for producing short-run TAB does not require the significant NRE of high-volume approaches.

Unistructure was selected because it provided the most cost-effective solution to the TAB supply problem. ISI also wanted to build an understanding of this process into the CAD environment to facilitate further experimentation and product development using this approach. Unistructure Inc. supplied design rules and tab frame design examples. The mechanical layout of the die bonding pads, die outline, TAB mounting frame, and outer-lead bonding pad locations were provided to ISI in the form of IGES files. Each of the IGES files were loaded onto Versacad at ISI to produce the mechanical TAB designs.

Two SRAM die types were required for the Encore memory module, the Hitachi HM6708 and Cypress CY7C192, meaning that two TAB frames had to be designed. Both memories

are high-speed 64K x 4 devices, but differ in that the '192 has 28 pads for separate I/O and the 6708 has 24 pads for common I/O.

On each device extra locations were provided to facilitate double-bonding of power leads. These extra bonds increased the lead count for the 6708 from 24 to 28, and for the 192 from 28 to 30.

The flip-TAB design provided face-down die attach with 0.030" leads. This allows a device placement pitch as small as 0.050." if footprint bounding boxes are allowed to overlap. A number of iterations were designed and reviewed by Unistructure to insure that the design met with the fabrication process requirements. The outer edge of the outer lead-bond (OLB) is .0375 inches from the die edge boundary.

The minimum inner-lead bond (ILB) was designed to fit on the industry-standard die bonding pad size of 0.004" square. The typical OLB pad was .005" x .015". A symmetrical and balanced design is required to prevent any lead distortion during the manufacturing process.

A .010"-wide kapton ring to maintain mechanical alignment of the leads was placed .0025" away from the edge of the die. Since the minimum OLB pitch of the TAB footprint was .008" and the memory bonding pads were placed on irregular intervals as small as .007", any misalignment between the ILB and the OLB was adjusted by angling the lead across the kapton to the next 8-mil increment. It was determined that if the maximum offset between the ILB and the OLB is within .0005 inch then the kapton ring or bar strip would not be necessary to keep the leads in proper alignment. The kapton ring was retained in these designs to support the die during the lead form process step after the leads had been excised.

The TAB leads change width as they progress from the ILB to the OLB. The lead at the ILB was .002" wide and maintained this width to the kapton support ring. The lead width then changes to .003" at the kapton ring and maintains that width out to the OLB. The lead narrows to .001" for .010" beyond the OLB; the inner "shoulder" of this narrow area is the target cutting line for the die excise. The lead width resumes it's .003" width to fan out to the JEDEC standard test pads at the periphery of the TAB mounting ring.

When the leads are cut they are formed with a .010" diameter arch, or "service loop," that is .003" to .005" high. The arch is used to minimize stresses caused by expansion mismatches between the die, the substrate, and the TAB lead itself. Since the TAB is mounted face-down, the only vertical offset required in the TAB lead is .003" to accommodate the die-attach- film.

After the initial exchange of IGES data, later designs were transferred to Unistructure using the DXF file format developed for mechanical CAD. This was requested by the engineering staff at Unistructure; in their experience DXF more accurate and complete than IGES.

In parallel with the TAB frame design, ISI initiated the design and fabrication of excise and form tools for these die configurations. ISI contacted two precision tooling companies to

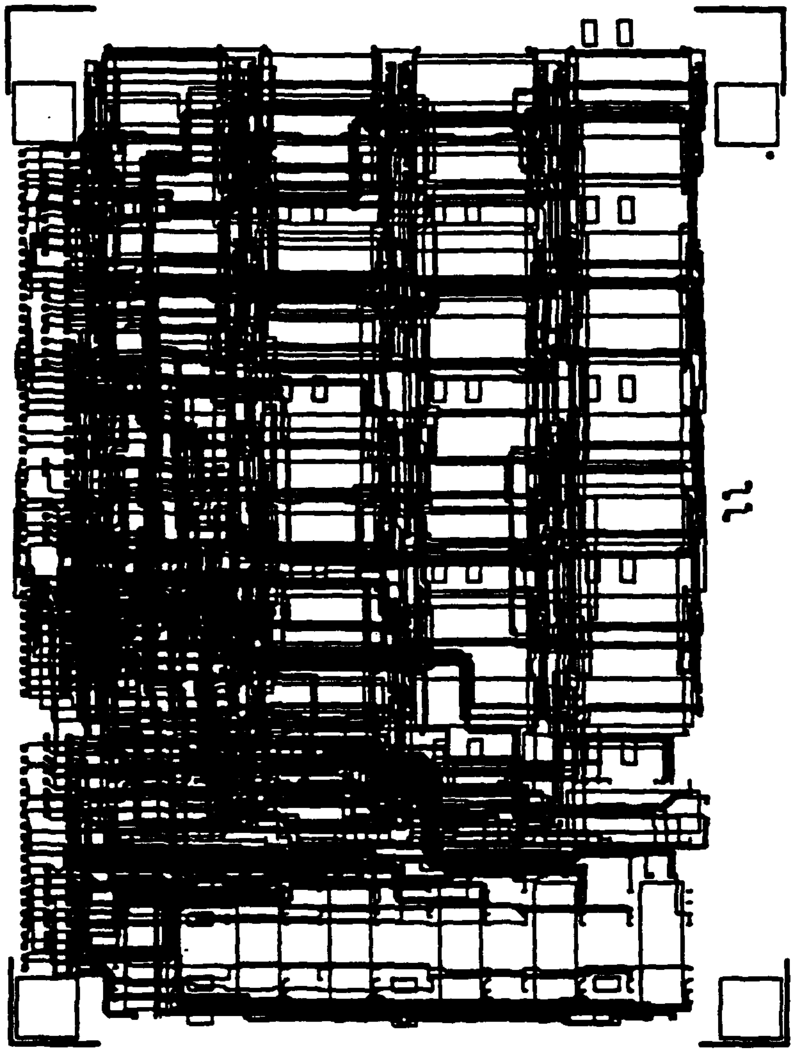

Figure 17: Encore Genesis Memory - TAB version (2X scale)

inquire about cost and lead times for supplying the custom tooling manufactured to ISI specifications. Detailed drawings modified from Unistructure's IGES design files were supplied to the tooling companies. Unfortunately, the lead times were too long to meet project deadlines so a decision was made to initiate an in-house, improvised tool design.

A substrate design was performed to mate with the TAB designs. The results of this effort are shown in **Figure 17**. In the TAB substrate, the primary thermal path is through the substrate. As a result the routing of the substrate appears much more clustered and denser than the routing of the SFLC substrate. This clustering is due to thermal conduction columns that do not appear in the figure.

TAB BONDING AND ASSEMBLY

The ILB attach process was a gold/gold or gold/tin weld or braze. It has been demonstrated that the relatively small TAB lead is not strong enough to apply significant loading to the

weld joint; embrittlement concerns developing from process metallurgy in large leads are not relevant to this approach. The ILB of the TAB-mounted die was overcoated with polyimide material.

After overcoat curing the dice were to be functionally screened. High-power devices must have the overcoat removed from the interior surface of the die so that thermally-conductive die attach material can be used for conducting heat to the substrate.

Tested devices were to be mounted upside-down on the substrate and lead-bonded using a welding technique similar to that used for the ILB. The OLB pad was approximately .005" x .015"; the leads of the TAB was approximately .030" long extending .020" beyond the die edge. One re-work cycle is supported with .030" leads and three re-work cycles could be supported with .035" leads. Re-work would be accomplished by pulling the die off the substrate, breaking the TAB lead at the weld joint. Elongated substrate pads allow repeated attachment of replacement parts by moving inward from the last bond location. Stress relief is built into the leads by forming the leads with a .003" arch or "service loop."

DIE SCREENING AND MODULE TESTING

A test procedure was created for testing memory die and an assembled HDSM memory module using an Integrated Measurement Systems (IMS) Logic Master tester. Test programs for the Cypress CY7C192 and the Hitachi HM6708 memory chips were prepared for the IMS. Since the memory devices are 64K deep while the IMS vector depth is only 16K, an external counter was required to generate device addresses under test control to relieve the overall testing time as well as test set complexity.

ISI designed an adaptor board for the IMS that included the needed address generator, a socket to accommodate devices mounted in a standard DIP packages, and connector footprint to accept the completed HDSM memory module. The design of this adaptor followed closely the development of the memory module, being designed and laid out in three weeks.

The two assembled memory modules required functional testing before shipment to Encore. A suite of memory tests, built from past experience and descriptions in the literature, was written. The tests include writing patterns affecting single bits (to check for shorts between bits), 4-bit groups (check memories at the package level), alternating-bit patterns ("checkerboard"), and a address-to-data test. The memory module delivered to Encore passed the bit-level tests within the signal and timing resolution allowed by the IMS. The more complex tests yielded mixed results conflicting with the results of the bit-level tests. The source of this conflict is believed to be a result of the adaptor design and is under investigation.

Testing memory devices is more difficult than testing logic. Since every memory location must be exercised in a variety of ways to detect both device bit-level errors and incorrect wiring at the module level, tests become brute-force pattern-generation and checking exercises. Experience with using the IMS on the Genesis memory has shown that although the external counter helps greatly in filling the memory with patterns, numerous small tests

must be written to overcome the lack of acquisition vector depth for reading data. A more sophisticated adaptor, or possibly an entirely new test system may be needed to handle module-level memory testing.

DESIGN TOOL SOFTWARE

The design of the Genesis memory module required use of a post-processing suite of programs written to extend our computer-aided design system, Omnicards. These programs read the design database generated by the Omnicards package and produce a mask for each process step, particularly those for vertical structures not encountered in printed-circuit processes. The post-processing tools were revised to accommodate changes made by Task Technology in their design database file format.

In addition to finding and generating masks for via structures, the program was modified to use an external description of vertical features. Normally, vertical structures penetrate a substrate only to the lowest layer where they are used. In the case of the memory module, the process requirements of the substrate fabricator required modifications to the existing programs; e.g., the columns of copper for connector pads had to descend at least to one layer above the metal backing the substrate so that the columns could withstand the force of the connector pins without allowing plastic flow in the polyimide.

A program used to generate plots of Gerber photoplotter files was extensively modified. Changes include determination of bounding boxes (scaling the pen plot using the largest and smallest X and Y coordinates appearing in a Gerber file), reporting the photoplotting apertures used, and reporting how many times the individual apertures were referenced.

DEMISE OF UNISTRUCTURE

The TAB and substrate designs were completed on December 19, 1990 and taken to Unistructure for review. During the review, the management of Unistructure was notified by their funding organization that continued failure to show either profits or a schedule for full-scale production warranted a total shut-down of operations. All projects currently in fabrication were lost. The fabrication processes and assembly details developed by Unistructure were dissipated with the relocation of the development staff.

MODULE INTERCONNECT

The air-cooled HDSM format uses a high-density, fine-pitch, stackable connector between the multiple layers of multichip modules to form a high-density module stack. The ISI-designed stackable connector uses button/plunger technology, a variant on the original button technology previously developed by TRW. This connector, (see **Figure 18**), manufactured for ISI by Cinch Connectors, is 3.5 inches long by .25 inches wide by .24 inches high. It has 222 contacts on .040 inch centers. To reduce crosstalk on large stacks, the connector has a controlled impedance of 63 ohms on the two outside rows of contacts. The center row of contacts has an impedance of 100 ohms, and can be used for power, ground, or other signals that do not have high clock rates or fast rise and fall times.

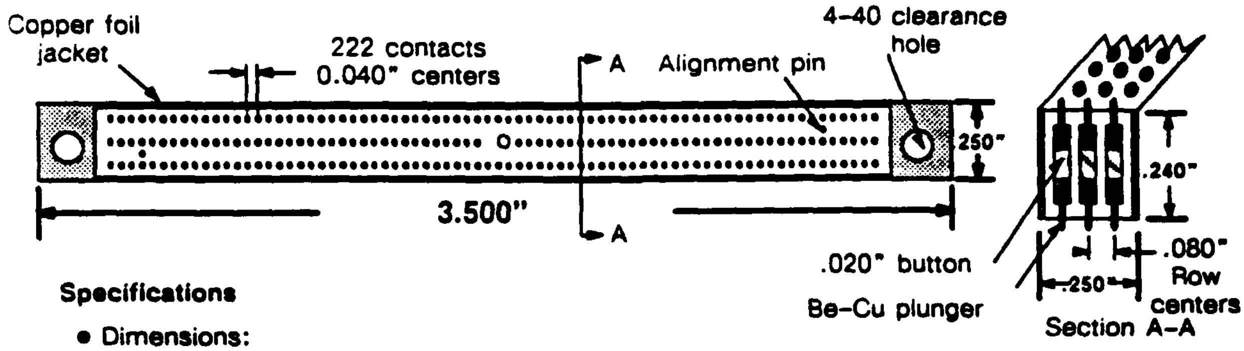

Specifications

- Dimensions:
  - 3.500" x 0.250" x 0.240"
- Features:
  - 222 gold-plated contacts on 0.040" centers
  - 312 contacts/in^2
  - 2.8 ounces/contact force
  - 63 Ohms impedance on outer rows
  - Commercially available
- Description:
  - A stackable connector for interconnecting MCMs in high-density systems

Figure 18: HDSM stacking connector

CONNECTOR QUALIFICATION

An incoming inspection test was performed on connectors received from Cinch. The test involved monitoring contact resistance while mechanically cycling the connector. To facilitate this test, two circuit boards were designed, built, and attached to either side of a connector under test. Traces on the board were designed to connect every pin on the connector in series. Each connector was given a serial number and tested through five cycles of compression or until the connector failed. Two additional cycles were done on several connectors that exhibited unusual behavior during the first five cycles.

Initially, only twenty-four of the initial 44 connectors provided by Cinch passed the incoming compression cycle test. Failure modes seem to indicate sticky pins in certain repetitive locations possibly indicating tolerance problems in the mold. Connectors that failed the initial test were returned to Cinch for analysis. Design changes in the connector were completed at Cinch and 10 prototype connectors of the modified design were fabricated. These new version connectors were tested by the same incoming compression test used for the previous connector version. The new connectors passed the compression tests with total resistance measurements for 444 contacts connected in series of 6.5-7.0 ohms.

After these initial compression tests were complete, an entire HDSM module mock-up with three layers and six connectors was constructed and a suite of environmental tests was run at Cinch. These tests included:

1. *Vibration per MIL-STD-1344, Method 2005, Condition IV.*
   This standard requires 20 g's peak for 4 hours in each plane while frequency is swept from 10 to 2000 and back to 10 hertz every 20 minutes.
2. *Shock per MIL-STD-1344, Method 2004, Condition E.*
   This standard requires a 50 g sawtooth shock of 11 millisecond duration in each plane.
3. *Shock per MIL-STD-1344, Method 2004, Condition C.*
   This standard requires a 100 g half sine wave shock of 6 millisecond duration in each plane.
4. *Compression cycles*
   This test involved a 1000 cycle compression test of the entire module stack by applying 75 pounds of force from a crosshead re-applied at a rate of 200 strokes per hour.

The fixturing for the shock and vibration tests included six (6) printed circuit boards furnished by ISI separated by two (2) Cin:Apse 222 position connectors per layer. This module was mounted on a 1/4" thick aluminum plate using 4-40 threaded posts at the ends and a 2-56 threaded rod at the middle. The module and plate were then fastened to a 1" thick aluminum plate which is used for mounting to the shaker and shock equipment with 1/4-20 bolts. The equipment used was: M.B. Electronics Model #N214 for vibration, Avco Corp. Model #SM105 for shock, and a detector built by Cinch to detect 1 microsecond discontinuities at 100 milliamps.

The results of testing are outlined below:

1. *Vibration testing:*

   a. No discontinuities in the plane of the contacts.

   b. No discontinuities with the long sides in a vertical position until at the 2-1/2 hour point a 2-56 rod snapped while passing through resonance. Discontinuities were detected in the two top connectors on the side with the failed rod.

   c. Without repair there were no further discontinuities with the long side in a horizontal position. In fact only the top connector showed a discontinuity in this position.

2. *Shock Testing*

   Other than the unsupported top connector, no discontinuities were detected in either the 50 g sawtooth or 100 g half sine wave tests.

3. *Compression tests*

   No failures were detected in 1000 cycles.

The results of these tests were very promising and the order for additional connectors has been released to Cinch. The mounting system needs improvement if 20 g variable frequency

vibrations are a possibility. This could be accomplished by a redesign which used three 4–40 threaded posts or possibly by increasing the strength of the clamping bars and using only the present two 4–40 threaded post. These possible modifications were deferred until a customer need arises.

Connectors were delivered to Berkeley, Harris, and MCC for further evaluation.

ENCORE GENESIS STATUS

A complete Processor module has been shipped to Encore and, according to an Encore memo, "..passed all diagnostic internal tests on the stand-alone bench tester and in a Multimax system." This Collaborative Development Effort is now complete.

MCC - ES-KIT / GENESIS

A collaborative project was undertaken by ISI and MCC to develop an ES-Kit-format board for demonstrating the Genesis processor module. The design and assembly of the adaptor was performed by MCC, the processor demonstration software was written by ISI.

A project was defined that included re-design of the standard ES-Kit 88000 processor board to accommodate two of the ISI-designed high-density dual-88000 processor modules. The re-designed board was intended to operate in the ES-Kit environment, however, the on-board facilities (including EPROM, RAM, and serial communications) allow stand-alone operation outside the ES-Kit to accommodate demonstrations at ISI. Design and fabrication of the adaptor board were performed by MCC.

Modifications to MCC's EEPROM-resident power-on diagnostics and rudimentary monitor were made by an APT systems programmer. The assembly was debugged at MCC and returned to ISI for demonstration.

The program chosen to demonstrate the multiple-processor module was the Sieve of Eratosthenes, a prime number generator. The Sieve program was first written and debugged on a SUN3 workstation. The program was converted to C++, the native programming language of the ES-Kit system, and cross-compiled for the 88k CPU. The demonstration program runs on SUN3 and SUN4 machines, as well as the multiprocessor Genesis/Es-Kit adaptor card.

The demonstration version of the Sieve used all available RAM except for per-CPU reserved areas. The computation effort could be partitioned among subsets of the four CPUs. Differences in measured elapsed time is displayed to show the effect of employing multiple processors. As shown in **Table A**, two CPUs speed up the Sieve by 80–113 per cent, depending on the range of numbers being searched.

Starting number		Elapsed time (seconds)		2-CPUs	
hex	decimal	1 CPU	2 CPUs	sigma	Speed-up
1	1	5.4536	2.5570	0.0069	2.1328
100001	1048577	5.6639	2.6713	0.0166	2.1203
1000001	16777217	6.4366	3.0815	0.0055	2.0888
10000001	268435457	7.7540	3.8825	0.0059	1.9972
20000001	536870913	8.2293	4.2041	0.0061	1.9574
40000001	1073741825	8.8015	4.6047	0.0053	1.9114
80000001	2147483649	9.5261	5.1296	0.0059	1.8571
FFE80001	4293394433	10.4521	5.8149	0.0067	1.7975

Table A. Multiprocessor Sieve program performance

The two-CPU case was run 40 times to get a reasonable statistical sample. The one-CPU case was run only 20 times because these latter execution times were quite consistent, never varying by more than 0.2ms between the slowest and the fastest times. The speed-up number is the ratio of the one-CPU execution time to the two-CPU execution time.

The better-than-2x improvement is attributed entirely to data caching: the innermost loop of the demonstration program is less than 1KB in size and lies entirely in one 4KB page. Therefore it should execute entirely out of the instruction cache.

When two or more CPUs were used, each CPU was given exclusive responsibility for sieving one part of the total memory; there was no need to maintain cache coherency between the 88200 data CMMUs. The only M-bus activity comes from filling cache lines and/or writing updated lines back to memory.

It was conjectured that, in the early stages of sieving when the multiples of smaller primes are masked off, there was enough spatial locality of reference to yield a high percentage of cache hits. The cache-hit percentage decreases when multiples of the larger primes are masked off.

The demonstration had 92KB, or 94,208 bytes, to use for its bit string representing odd-numbered integers to be sieved; thus, each invocation found any and all possible primes in a range of 1.507,326 numbers. The starting numbers were chosen at random in an attempt to understand program execution when differing numbers of passes through the bit string were required. The numbers chosen were not special except the last one, 0xFFE80001, which is probably the largest initial number that will yield a range of primes expressible as unsigned 32-bit integers.

Starting number		# primes	last prime	# primes	density
hex	decimal	sieved	sieved	found	
1	1	197	1229	114700	0.0761
100001	1048577	248	1601	104841	0.0696
1000001	16777217	584	4283	90473	0.0600
10000001	268435457	1901	16433	77573	0.0515
20000001	536870913	2587	23209	74981	0.0497
40000001	1073741825	3513	32797	72418	0.0480
80000001	2147483649	4791	46381	70187	0.0466
FFE80001	4293394433	6539	65537	67814	0.0450

Table B: Multiprocessor Sieve program statistics

Density is the fraction of primes found in the range of numbers sieved. In the case of this trial, the number of primes found divided by 1,507,328. Statistics developed from the Sieve program running on the Genesis / ES-Kit are presented in **Table B**.

Sun-3 and Sun-4 implementations were run to verify of the 88000 implementation of the Sieve of Eratosthenes. Identical results were produced.

INTEGRATED SILICON MICROPHONE

INTRODUCTION

APT used a VLSI silicon microphone element developed at UC Berkeley as the basis for a hybrid demonstration effort. The demonstration used hybrid packaging technology to rapidly produce a small-scale system prototype. The experiment demonstrated not only the microphone element itself, but also the LagerIV VLSI design system, standard cell support from ITD, custom signal-processing chip designs from UC Berkeley and UCLA, custom chip fabrication by MOSIS, and packaging technology and system integration from APT.

PACKAGING APPROACH

The Microphone project demonstrated packaging approaches at Level I, Level II, and Level III. The Level I approach mounted dice directly on a custom designed, 1-inch-square, multilayer ceramic substrate. The die-attach was conductive epoxy with die interconnected via wire bonding. The substrate, or Level II package, contained about 25 surface-mount parts including the microphone die, a custom VLSI signal-processing die, a crystal oscillator, a set of operational amplifiers with gain-selecting resistors, an 8-channel analog-to-digital converter, and numerous discrete devices (see **Figure 19**). The substrate was "programmable" in the sense that each functional block on the substrate had jumper options to allow modification by wire-bonding at final test time. This approach demonstrated the concept of programmable interconnect used on "standard" modules.

PROJECT DESCRIPTION

A multi-layer, thick-film hybrid substrate was designed and packaged in a custom Kovar hybrid package. The one-inch-square substrate supports the .2 inch square microphone element, a commercial amplifier chain, a commercial A/D die, a crystal oscillator, and a custom signal-processing chip. The substrate and microphone are aligned over an acoustic port drilled in the package. The package lid forms the sealed chamber for proper operation of the microphone. The digital signal processing chip detects acoustic energy in a narrow frequency band around 2 kilohertz.

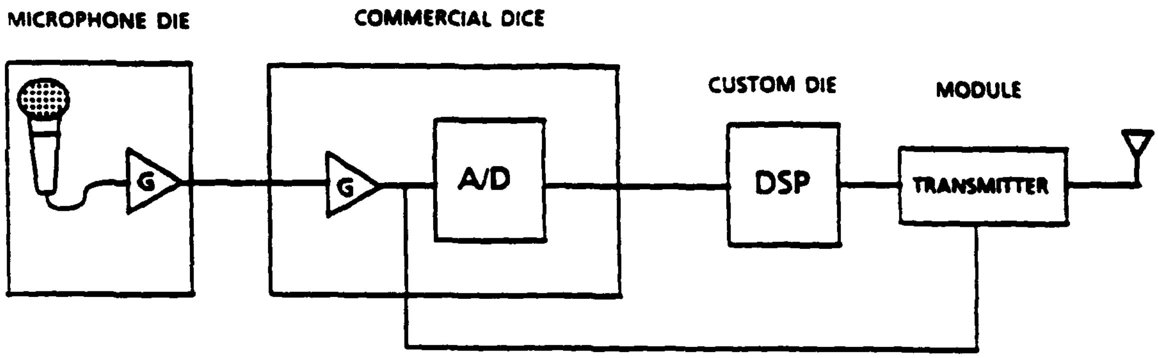

Figure 19: Microphone Block Diagram

The Level III package consists of the substrate mounted in a standard hybrid flat-pack with radial leads. This approach allowed the microphone element and associated electronics to be sealed to MIL-SPEC-883, except for one static pressure equalization port.

The hybrid package (see **Figure 20**), which measures approximately 1.2" x 1.2" x .2", was mounted in a custom-machined, two-part plastic enclosure with batteries and a simple radio transmitter. This enclosure was intended to demonstrate the ability to machine a plastic part that can be used as a pattern for molding additional units. The enclosure itself is a demonstration of "standard frames" in that it contains several "shelves" that could be used for additional hybrid packages to provide additional signal processing, data storage, or communication support.

DESIGN METHODOLOGY

The overall system design methodology is to create the system components as "standard frames" so that they may be assembled in application-specific ways at system deployment time. In addition to the requirement to build a compatible set of system components, there

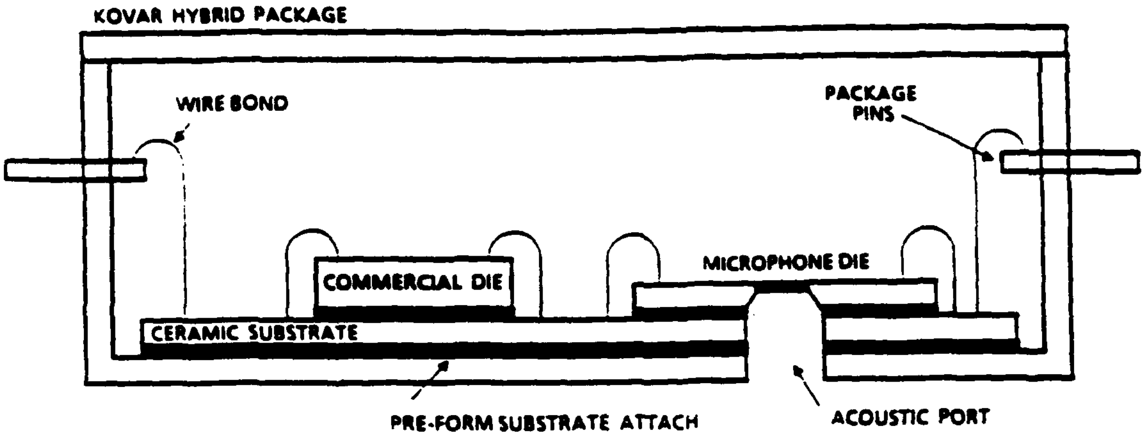

Figure 20: Microphone Hybrid Package

are several other interesting and unique problems encountered in this type of system design. Chip designers, for example, must be concerned about low-power design approaches and supply voltages that can vary by a factor of two. Specifically, this project forces "real-world" system-level design considerations on chip designers. An additional application-specific design constraint in the form of testability is encountered because of the requirement to screen dice before they are mounted on the substrate. VLSI designers are free to determine an approach to die test, but the test must screen the entire functionality of the chip and may require no more than 8 I/O pads, including power and ground. This restriction allows several possible silicon standard frames to be tested on a conventional probe station with individual probes, avoiding the costs involved with separate design-specific probe cards.

ACCOMPLISHMENTS

Five hybrid substrates were pre-assembled and tested at ISI without microphone elements. The most critical component was the commercial amplifier, configured as a unity gain buffer for the microphone element. Input impedance of the buffer amplifier was 10^{12} ohms, and the input-referenced noise at 100 hertz was 10 nanovolts. The high input impedance of the buffer amplifier made it susceptible to induced noise and to offset drift that was caused by a charge build-up on the capacitance of the microphone element. A reversed-biased, surface-mount diode was used at the amplifier input to.to remove the built-up charge. During early stages of testing, however, the buffer amplifier and additional gain stages demonstrated significant operating point instability caused by the induced charge. A DC feedback loop was designed and constructed to stabilize the amplifiers. The A/D die and the UCLA-designed custom DSP chip were operational. 3.5 volts peak-to-peak of analog input signal to the A/D chip at 2 kilohertz were required to trigger the DSP energy detector. This level

was above the expected threshold of 2 volts peak-to-peak because of a DC threshold shift in the input signal to the A/D that was not anticipated in the DSP chip design.

The sensitivity of the microphone elements was disappointing, as it proved to be about 1/250 that of a normal microphone. The sensitivity level of the two microphones was measured at four kilohertz with an "A" message-weighted filter. The results were 4.24 microvolts per microbar and .74 microvolts per microbar. These numbers compare to 1000 microvolts per microbar for a conventional electret microphone. Extreme sensitivity to incident light on the microphone diaphragm was also observed. Considerable 60 hertz noise was induced by nearby incandescent lighting when that light was allowed to reflect into the acoustic port. Because of the high noise component, about 70 dB, measurements were made with incident acoustic levels around 100 dB. In the final analysis, however, it was not high noise levels but insensitivity of the element that limited the usefulness of the tested microphones.

The reduced sensitivity of these microphone elements is caused by deformation of the acoustic membrane during the CMOS fabrication process steps. This deformation in turn causes the microphone membrane to develop stresses that result in the low sensitivity. Previous microphone fabrication runs without the CMOS process steps resulted in better sensitivities. UC Berkeley is proposing process changes that will alleviate stresses in the microphone elements.

UC BERKELEY / USC - BAM

The Berkeley Abstract Machine (BAM) project, originated at the University of California, Berkeley, has developed a machine architecture optimized for PROLOG. Moved to USC (and renamed Aquarius III), the BAM project developed a single-processor, SUN workstation-based evaluation board, BBGUN, to support the custom VLSI BAM processor chip. This board houses the BAM processor in a 299-pin PGA package, high-speed instruction and data cache memories, VME-bus interface, and random support logic totaling around 210 devices.

TASK DEFINITION

Discussions with UC Berkeley and the University of Southern California focused on desired architectural and performance goals for Aquarius III. These discussions resulted in a working document, which served as a specification for the implementation effort. Preliminary results of these discussions indicated that APT would assist in the construction of a single node prototype that would plug directly into an existing commercial workstation. The intent of this preliminary effort was to provide a hardware platform for debugging the architecture, developing the software, and supporting specific additional system designs such as the high-speed busses that were used to interconnect multiple nodes.

AQUARIUS III

The design transfer process is a critical aspect of a proposed Systems Assembly Service program. As part of an early investigation of high-level design information, ISI requested that the BBGUN design be made a candidate for a design transfer experiment. After agreement on goals, the effort was launched.

The design was transferred as a ViewLogic database. Using the schematic as the transfer medium allowed ISI to directly perform engineering consulting services for the board designers, reducing the overall time for design completion. Design netlist and partlist data was extracted from the database using ViewLogic tools and transferred to the PCB layout system for design rule checking.

The initial result of the incoming DRC was a large number of data syntactic errors that would be rejected by a service. Several iterations of design submission, incoming DRC, and design modification were required to develop an acceptable design.

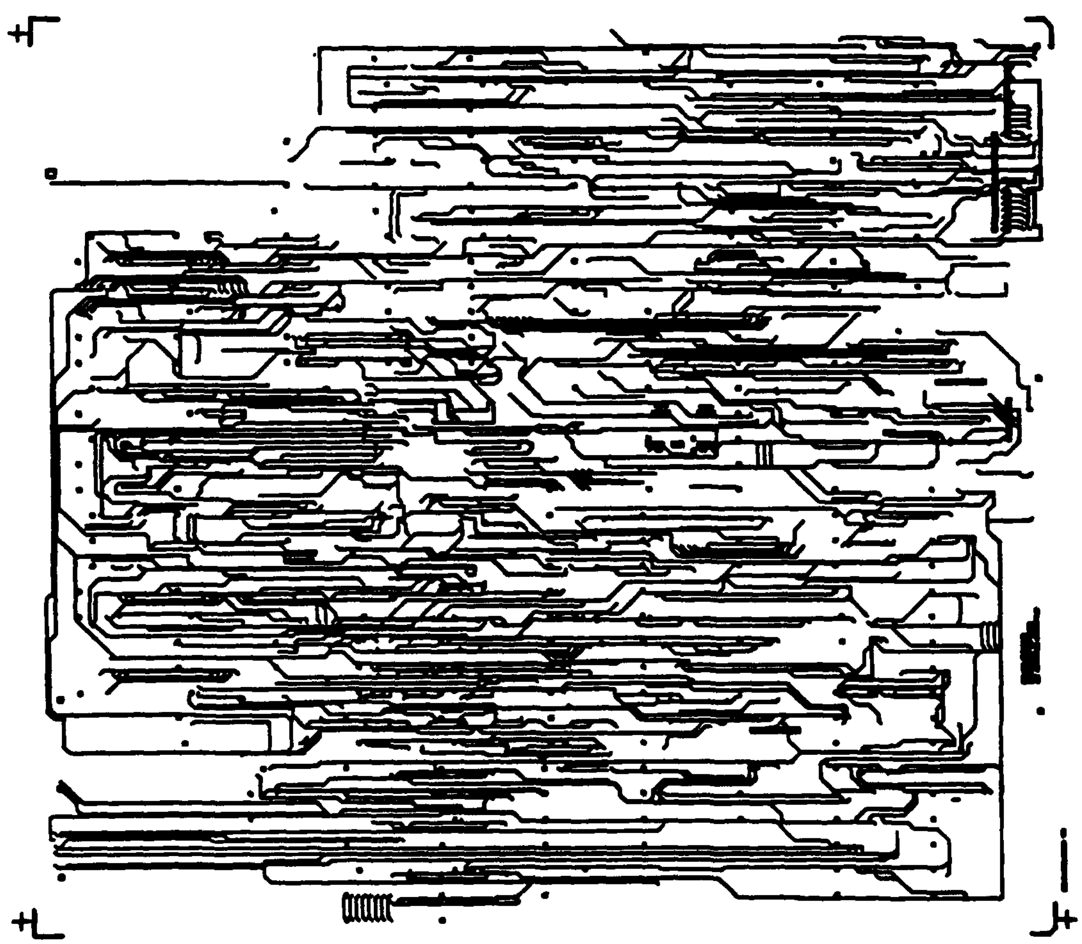

Figure 21: BBGUN Processor Board (top layer)

As an experiment, without an in-depth knowledge of the circuit, the components for the board were placed and a baseline wire routing performed. To verify that a reasonable part placement could be performed without *a priori* knowledge of the circuit topology, ISI requested a sample placement to compare against the baseline layout. The layout complied with the USC-provided sample to approximately the 95% level. After making changes in device placement, the board was re-routed and the new design returned to USC for review.

That only 90 minutes were required for these complete routings suggested that a more aggressive, reduced layer-count design could be implemented for essentially the cost of processing time alone. The number of routing layers was accordingly reduced from 6 to 4 and the router restarted. A complete routing of the circuit was produced in 6 hours, evidence that technology usage rules employed by a *Systems Assembly Service* form a complex decision space, with design complexity, processing time, and fabrication costs becoming factors in implementation decisions.

The BBGUN board was fabricated with an 8-layer stack in an attempt to produce low signalling noise levels. The results of the layout effort are shown in **Figure 21**.

FAST was used as the vendor for the BBGUN board assembly components. With some changes to local procedures, FAST effectively reduced the time required to administer the part procurement process. Methods for automatic component part list submission to FAST is under study.

Outside services were arranged for board assembly. Since the BBGUN is a first design for a new VLSI device, ISI recommended that every active device on the board be placed in a socket. Reducing the BBGUN assembly process to device socket and bypass capacitor insertion, followed by wave soldering and cleaning. The rule for recommending whether to socket all devices on a board is an area being studied for Systems Assembly.

The assembled board was returned from assembly and delivered to USC. The insertion of active devices was performed by USC project members according to their initial debugging procedures.

The BBGUN/BAM system has been shown to operate reliably at 30MHz. This operating frequency limitation was reached by the conventional packaging used for the ICs; the packages preclude denser packing required to reduce signalling delays caused by wiring length.

UC SANTA BARBARA SHUNT

SYSTEM DESCRIPTION

The SHUNT system is a 16-processor multicomputer with a connection-switched crossbar interconnection mesh. To keep system implementation costs within budgetary limits it was critical to facilitate fabrication and assembly. To this end, the crossbar interconnect was implemented with printed circuit cards; the interconnect was partitioned into a 9Ux400

VME-format backplane and a single high-density daughter-card. A schematic view of this arrangement is shown in **Figure** 22.

The backplane card is a 20-slot, 4-layer printed circuit card that provides signal and power interconnect for the SHUNT switch, the 16 custom processor cards, 2 SUN VME host processor cards, a system control processor, and sites for jumper cables from the top edge of the system controller. That is, the VMEbus signals appear only on backplane slots 1-3, the remainder devoted to the special interconnect required by the 16 processor cards. The backplane card layout, which uses relatively conventional 8-mil design rules, was completed.

The SHUNT switch was implemented on a 2-sided surface-mount board. This 8-layer printed circuit uses fine-geometry design rules (5-mil lines, 0.014" vias) to implement the wiring of the crossbar. Crossbar interconnect provides a pathological test case for most wire autorouters; the first successful routing of the switch took 70 hours to complete on a SparcStation 2.

Z-AXIS CONNECTOR

The daughter-card would be attached to the backplane with six 210-contact "PAI" z-axis connectors (see **Figure 23**) fabricated by Augat. These connectors are held captive to the backplane using conventional soldered through-hole pins on one side. The other side of the PAIC connector uses a surface-contact spring-loaded plunger to provide a high-normal-

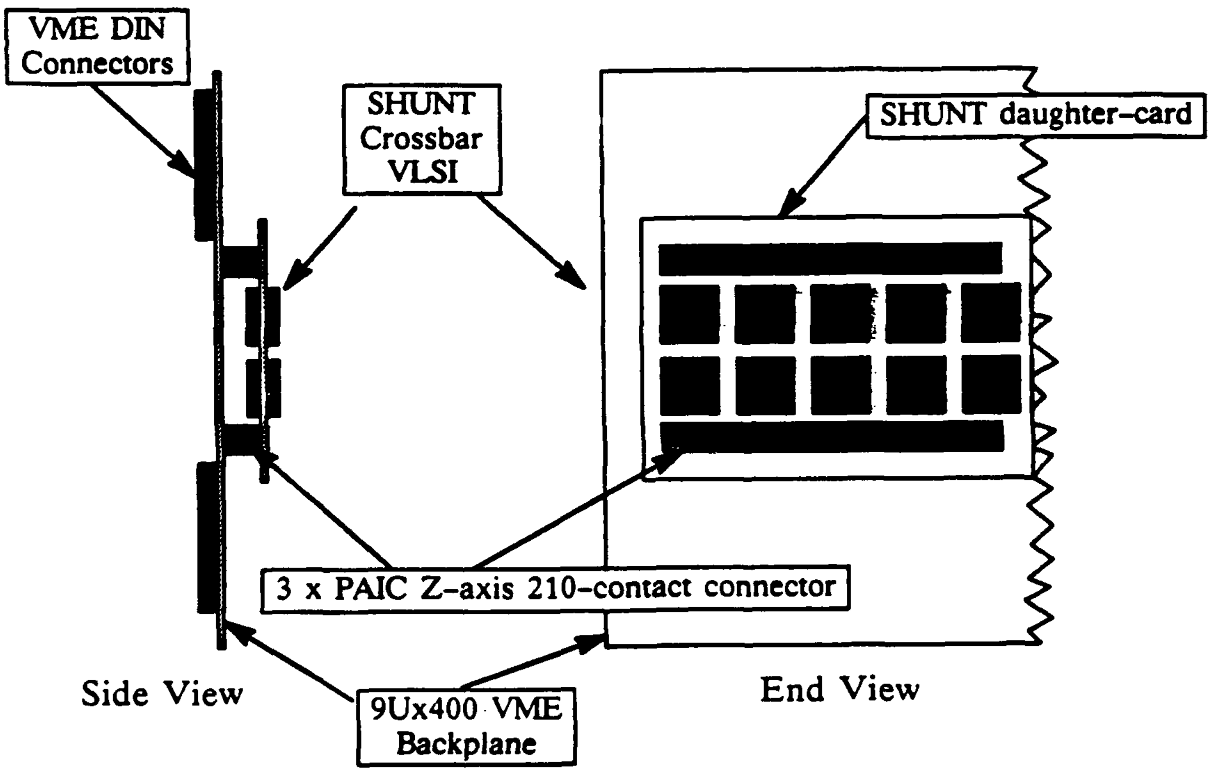

FIGURE 22. SHUNT System Backplane

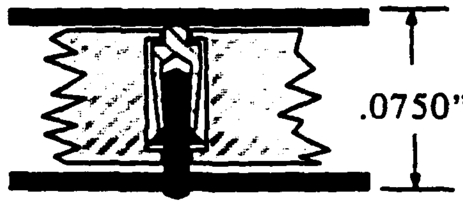

Figure 23. Augat Through-hole PAI Contact

force contact with the daughter-card. It was presumed that high area pressure exerted by this contact should provide reliable connections that could be easily removed for maintenance and repair of the daughter-card.

DAUGHTER-CARD COOLING

UCSB simulations indicated that the daughter-card VLSI devices would dissipate a maximum of 0.1 Watt per device. At this low power level it was concluded that, convection should be adequate to maintain a 60°C device operating temperature, even for those devices mounted on the daughter-card next to the backplane. If the device power levels had proven to be higher a small fan, mounted at one end of the backplane card cage, would have generate adequate airflow to maintain this operating point.

VLSI Functional Testing

SUNKIT IV TESTER

INTRODUCTION

Access to the DARPA-sponsored foundry service has provided the research community with a simple, uniform interface to fabrication of custom integrated circuits. As a result, the research community is faced with a need to test and verify a great diversity of devices. Performing this test and verification function has traditionally been costly, requiring significant capital expenditure for tester hardware.

The KITSERV project focused on research and development of functional tester systems intended to lower the cost of prototype VLSI device functional testing for the DARPA community. APT continued tester development begun by KITSERV, developing SUNKIT III, a low-cost, high-performance tester architecture. Based on the concepts of "event-driven" simulation and integrated test systems, SUNKIT III was targeted for technology transfer as a commercial product; to serve as a packaging demonstration project for APT; and to fill a functional test need at ISI.

The effort resulted in SUNKIT IV, a flexible functional tester architecture composed of custom VLSI, commercial memory products, and high-density, bipolar drive modules. It was to serve as a packaging demonstration for APT.

PROJECT DESCRIPTION

The SUNKIT IV prototype was to use a custom-packaged PinDriveIV VLSI device consisting of vector output and DUT sampling and error checking, timing-edge assignment, and accurate timing-edge placement. The PinDrive device develops the idea of "formatless" testing, allowing a use to change the data presented to the DUT on a vector-by-vector basis.

SUNKIT IV architecture is extensible to an arbitrary number of test channels. The VLSI PinDrive device was being designed to interface with a variety of memory devices, providing flexibility in implementing a demonstration system. Each PinDrive device was to contain timing-edge assignment and de-skewing hardware, moving the tester toward "per-pin" architecture while retaining a simple, low-cost implementation.

SUNKIT IV architecture provided a mechanism for performing high-speed wafer-probing experiments. Commercial test systems typically connect to wafer probe platforms via cables, degrading test system performance. In contrast, the entire SUNKIT IV package was smaller than many commercial probe station test heads, allowing the tester to be mounted directly onto the wafer probe station and reducing signal path lengths to three inches or less. High edge speeds intended for SUNKIT IV would allow high-performance wafer-level testing.

TEST GENERATION AND DISPLAY

ViewLogic, a commercial computer-aided engineering (CAE) system, was adopted as a test-generation front-end and display back=end. The requirements for such a front-end include flexibility over a variety of design philosophies and styles, generality in handling designs of arbitrary size and complexity, and a published, simple simulator interface.

TESTER SOFTWARE

The user software environment was integrated into the ViewLogic design environment. A software tool called *gen2sk*, which generates vectors from ViewLogic simulator output files, was written and demonstrated, allowing users to do physical design verification from a ViewLogic design and simulation database. Another tool, *skpost*, compared post-test results to simulation results. A source-level debugger was written to allow rapid tracing of test errors back to test set data. Test management software, *skmgt*, has been specified and is described in detail below.

Skmgt is an application to manage a SUNKIT IV tester. This window-oriented software was originally written to run under Sun's SunView system. **Figure 24** shows the major windows available to the user. The tool's main window details which test the user has selected to run and also contains buttons for performing operations as well as display area for status and error messages. One button calls up a window wherein one can "browse" through a directory containing several test modules, and select the desired one. Once the selected test module is loaded, other buttons (and windows) summarize the contents of that module and the

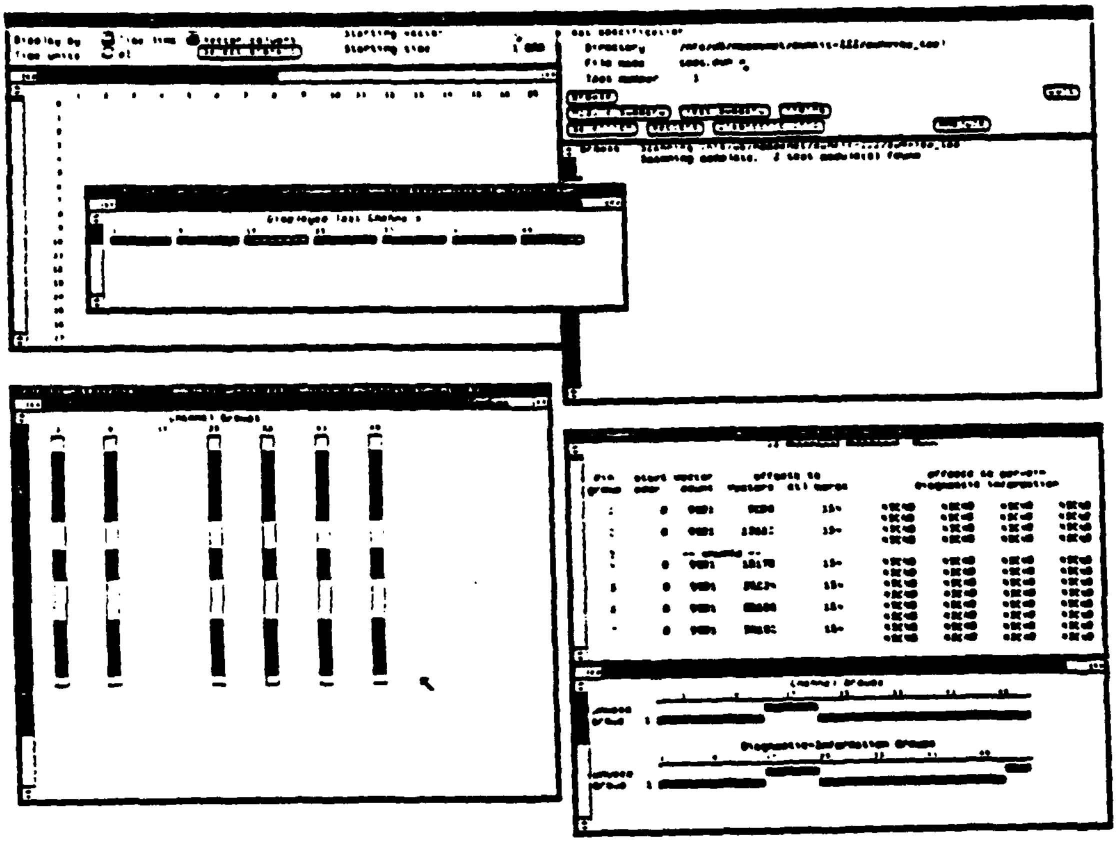

Figure 24: SUNKIT IV test management software, *skmgt*

selected test. Another button allows selection of vector ranges which can be displayed either in columnar form or in a manner similar to an oscilloscope. Another button allows display and alteration of timing for the test. Finally, there is a button to display the test results; again, the user can select the ranges of results vectors to be displayed, in either columnar or oscilloscope-like form.

APT investigated linking Viewlogic's circuit design and simulation suite of tools to the SUNKIT IV tester, and found that only indirect linkage through ASCII files is possible; there are no other "hooks". In other words, a ViewLogic component would generate an ASCII file containing, say, a series of vectors from a circuit simulation. The user would direct ViewLogic to invoke a program to run a test on the SUNKIT IV. This program would read the ASCII file from ViewLogic, generate SUNKIT IV vectors, load them, run the test, format the results into an ASCII file, then terminate. The appropriate ViewLogic component could then read the new file.

To this end, programs written for the SUNKIT II were extended to read ViewLogic Generic Waveform Files (the ASCII interchange files mentioned above), generate and load the resul-

tant vectors into the SUNKIT IV, conduct the test, and format the test results into another ViewLogic GWF file.

A Sieve compiler originally written for SUNKIT II, was modified to generate SUNKIT IV vectors. The Sieve compiler was also modified to provide a higher level test compiler for the low-cost CADIC tester used by ISL to debug SUNKIT IV test chips.

VLSI DEVELOPMENT

Considerable effort was focused on VLSI development activities. Among these activities were porting previously developed cells into the Lager environment, developing new cells, testing submitted VLSI devices, and designing new test devices for fabrication.

APT VLSI cells were added to the design environment library. Input capacitance and load capacitance derating factor for each of the leaf-cells were extracted and calculated with HSPICE. These values were added to ViewSim simulation models, allowing more accurate predictions of the performance of devices to be made.

Additional effort was made to design and characterize new VLSI leaf-cells. The sk_x1340 is a stackable tri-state buffer element for driving internal busses, which matches the width of an existing tiny latch family. The sk_xbena is a controller for up to 32 sk_x1340 buffers. Two cells were also constructed to allow up to three device probe pads to be placed in a standard cell array without disturbing power rails or wiring channels. Several more leaf-cells were built, including the sk_1580a master/slave negative-edge-triggered flipflop; sk_crs and sk_irs R/S flipflops, sk_crs2 and sk_irs2 R/S flip-flops with two reset lines; and sk_ipp & sk_cpp probe pad cells.

A test-circuit device submitted prior to the installation of the Lager design system was returned from fabrication. This device contained a complete single-channel timing generator. Several problems were found in this chip, but a large number of test probe points and laser cuts allowed most of the chip to be tested. We were able to repair and test the timing modulus counter using the laser. The fine-delay generator worked correctly. The coarse-delay unit was incorrectly wired between its control data latches and master control logic, preventing this feature from being tested.

HSPICE simulations performed on extracted geometry indicated that the latch structures should work correctly, showing that the tiny latch read-back circuitry could drive a 50MHz digital signal with a substantial capacitive load (440fF).

Two additional TinyChips were submitted for fabrication. One device contained two versions of tester output logic. The second TinyChip contained two versions of the tester acquisition logic. The core layouts for these devices were generated with Lager from ViewLogic schematics, with final routing to the pad frame performed by hand.

An experimental timing generator was laid out using Lager, followed by hand-editing connections from the circuit core to the pad frame. This device was mounted in a standard 4600u x 6800 microns MOSIS frame with APT custom I/O pads.

A complete PinDriveIII tester chip layout was begun, containing four complete sets of channel logic and a 20-bit address generator for vector sequencing. Included in the logic were channel-timing generators, and drive and acquisition circuits. The goal was to lay out this circuitry on a 6900x6900 micron die, allowing four devices to be placed on a MOSIS 1.6 micron reticule. After building a single timing generator with Lager, it was discovered that four timing generators were 7200 microns high, overflowing the planned die size. This problem with VLSI layout forced a re-evaluation of the goals of SUNKIT III and led to a re-partitioning of the design, creating SUNKIT IV.

TIMING GENERATOR

Fabricated in 2-micron CMOS, this device provided fine and coarse timing edge-placement and frequency pre-scaling. **Figure 25** shows a recording of 8 of the 32 possible fine-incre-

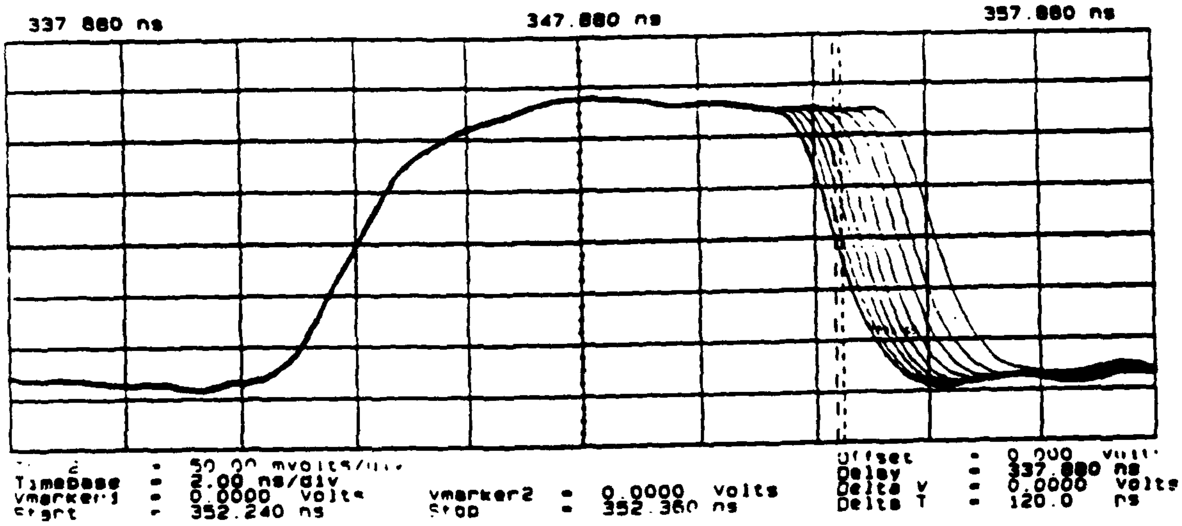

Figure 25: Timing-edge placement

ment edge-placement intervals. This plot demonstrates the ability to provide timing-edge placement in 100- picosecond (250 picosecond maximum) increments over a 2-nanosecond range. The coarse edge placement selects 1 of 16 fine-placement intervals, and the frequency prescaler selects 1 of 256 coarse-placement intervals. The result of this selection is that an edge can be placed to within 250 picoseconds anywhere within a 7.68-microsecond window. The frequency prescaler operates at frequencies up to 60 megahertz, or twice the required speed.

Stability and phase jitter between timing edges is at most 200 pS, including drift in the test system clock generator. **Figure 26** shows a record of phase jitter between timing-edge signals over a 5-second interval.

Fabricated in 2-micron CMOS, this device produced an output pulse whose width is equal to the time difference between two input pulse edges. **Figure 27** shows the minimum full-height pulse width that can be generated this circuit. This 3.3 nS wide pulse is three times the required performance.

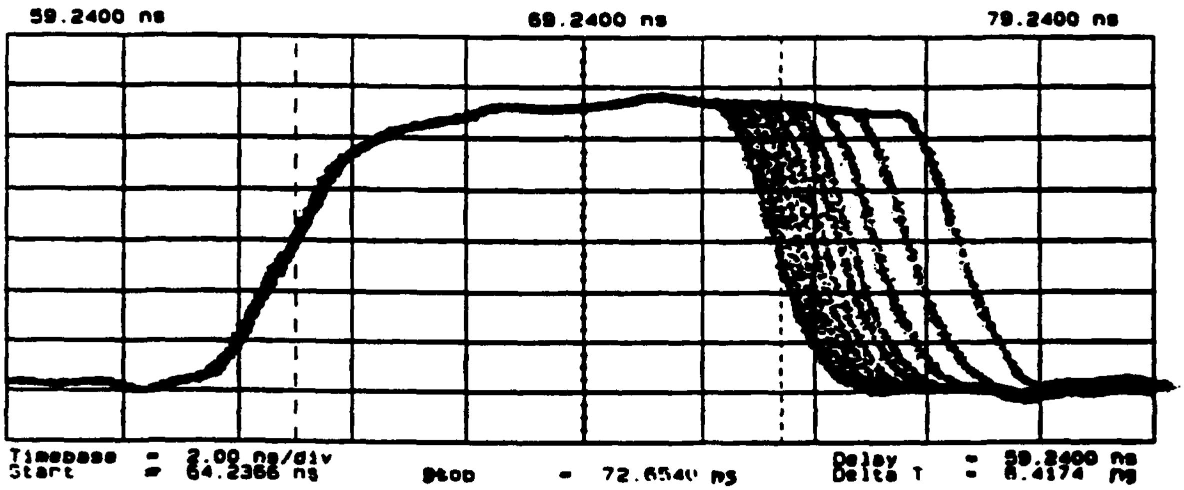

Figure 26: Timing-edge phase jitter

PACKAGING

A custom VLSI package was designed for the PinDrive device, 220-pad leadless surface-contact interconnect. The large I/O count allowed multiple high-speed I/O signals to be implemented on the device while providing a conservative signal-to-power-pin ratio. Pin grid arrays have parasitic capacitances and inductances starting at 5 pF and 100 nH. By comparison, calculations show that the PinDrive parasitic lead capacitance and inductance should be on the order of 1 pF and 2 nH. Reducing power rail inductance will greatly reduce device power noise, extending the range of device performance and eliminating a source of possible latch-up problems. Reducing signal capacitance simultaneously increases system speed and reduces I/O driver size.

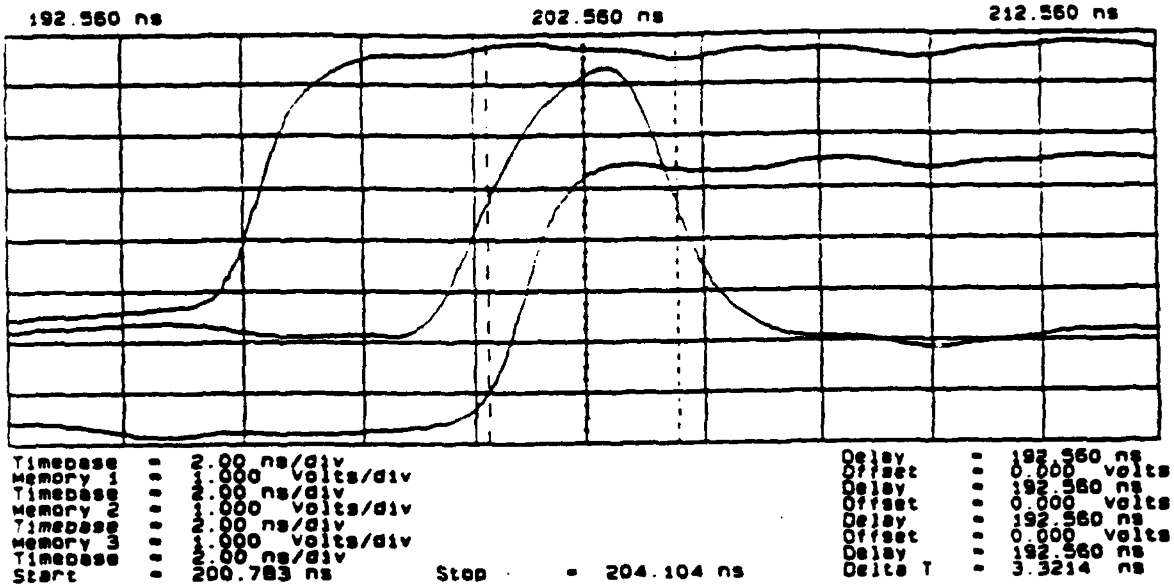

Figure 27: Minimum-width drive circuit pulse

Technology Development

INSTALLATION AND SUPPORT OF A CAD ENVIRONMENT FOR VLSI DESIGN

This report describes the consulting support provided to ISI for creating a VLSI design environment based on public domain tools compatible with the MOSIS supported CMOS processes.

This task not only involved installation and development of software but also an understanding of how the VLSI design task should be partitioned in a team and how the hand-off should occur between different team mates. This partitioning and hand-off is naturally affected by the CAD infrastructure.

BACKGROUND

The driver for this project was the design of a tester chip. At the start of the project, the design tools in use were ViewLogic for system design and Magic for layout. The system designer was creating schematics in ViewLogic and simulating with Viewsim and then handing over specifications of various modules to the layout designer. The problems encountered were a) the layouts did not match the system designers expectations, b) the layout designer did not have an overview of the entire chip. It was felt that both problems could be avoided if an integrated environment existed which allowed both system and layout designer to exchange design information via a CAD database.

CAD DEVELOPMENT

The above approach required a link between the ViewLogic schematic entry and simulation and the Magic layout environment. The use of the LagerIV design system for this purpose was explored. To enable the use of LagerIV a link had to be developed between the ViewLogic database known as Viewbase and the LagerIV database known as OCT. A database translator called vb2oct was written for this purpose.

Following this, an integrated environment was developed whereby the system designer specifies designs using ViewLogic to describe a schematic of standard cells available in the LagerIV library. The standard cell schematic is handed over to the layout designer. The layout designer generates the OCT database using vb2oct and then runs the LagerIV tools to generate the layout. He evaluates the layout for performance and functionality using Spice and IRSIM. The functionality can be checked by generating IRSIM test vectors from the test signals specified in Viewsim by the system designer. If the performance is not satisfactory, then new cells are designed or existing cells are modified.

STANDARD CELL SUPPORT

To effectively use LagerIV, a stable cell library and place and route tools were found essential. Several new cells were developed by ISL Staff as part of this effort for the standard cell library in LagerIV. Furthermore, this project provided extensive evaluation of the standard cell place and route capabilities within LagerIV and led to several improvements and debugging that were carried out with ITD and MSU.

MAIN RESULTS

While most of the requirements for the tester chip could be met with the integrated environment, one limitation was the lack of performance driven design tools. In the tester chip, in one module it was critical to balance the delays. This requires a careful placement of the standard cells. However, since the placement is automatic and could not be manually controlled, this module had to be laid out by hand.

The final result of this project was: a) the installation of LagerIV at ISI, the development of the vb2oct data base translator and, the training of the layout person at ISI to use the integrated environment. b) Extensive evaluation of the standard cell layout support in LagerIV. c) An understanding of what is an efficient approach to VLSI design.

These are described in detail below.

vb2oct DEVELOPMENT

The first step in creating vb2oct was to identify counterparts between the database objects in Viewbase and in OCT. This turned out to be feasible. The second step was to write code using the Viewbase support library to read Viewbase netlists and create corresponding netlists in OCT. The OCT view used is the structure_master view as defined for LagerIV. One

of the things we learned is that it is necessary to track upgrades in LagerIV which affect the OCT view definition. Other than that, vb2oct did not require much maintenance once it was developed.

In the initial stages considerable amount of debugging was required to make vb2oct work. This could have been avoided with more documentation on Viewbase. While the documentation is adequate for software development, troubleshooting requires some experience with the Viewbase library.

A second issue we had to deal with was maintaining consistency between the ViewLogic library and the LagerIV library. For each logic cell there is a ViewLogic symbol and simulation model. At the same time there is a corresponding layout cell in the LagerIV library. Changes in the cell layout need to be propagated to the ViewLogic library if they affect the I/O or logic function.

DESIGN HIERARCHY ISSUES

An interesting problem encountered was that the hierarchy used by the system designer in the ViewLogic schematic was not necessarily the best way to partition the layout. The hierarchy in the chip architecture is defined based on the functionality of different blocks and the ease of representing the design. For the layout efficiency, it was found necessary to flatten parts of the hierarchy and treat them as one composite circuit. To achieve this, the FLATTEN feature of LagerIV was explored. The database translator was modified to allow use of the FLATTEN feature. Significant reduction in chip area was observed by flattening.

In the course of the design several bugs were found with the standard cell place and route software in LagerIV. For example, a pathological problem was the appearance of stray metal lines in the layout. In all cases the layout designer successfully interfaced with support people at ITD and got bugs removed. Updated versions of the code were installed at ISI. Initially the layout person needed assistance with the code installation however by the end of the project was able to independently recompile and install the code.

DESIGN MANAGEMENT

Given the above design system the question is how can a design team work efficiently. One of the main issues is who is responsible for correctness of the schematic and who "owns" it. The model we experimented with is that the system designer owns the schematic and only he changes it. The layout designer has to be able to take the schematics and generate the layout from it. This implies that the layout designer has to learn ViewLogic. If the layout designer finds a better way to describe the schematic (because of his intimate knowledge of the cell library) he is not allowed to arbitrarily change it but changes it in consultation with the system designer. This was observed to be a bottleneck.

Another issue investigated is consistency between the schematic and layout. Further work is needed to address this issue. However, existing netlist comparison tools in LagerIV can be

modified to achieve this. A useful utility would be to generate IRSIM test vector files from Viewsim test vectors. That way identical test vectors can be used for the Viewsim simulation on the schematic and the IRSIM simulation on the layout. Another utility for automatic comparison of the output vector can then verify functional correctness.

A third issue is, who is responsible for the CAD system itself. The model we experimented with is that the layout person maintains the CAD system with support from MSU/ITD. This worked quite well and by the end of the project, the layout person was able to install upgrades to the software.

The recommended design management approach, based on the above experiments is that the layout person be put "in-charge" of the chip design and interface with the system designer at a very high level. At the start of the project the interface was at the level of Magic modules. This was clearly too low a level and with the work done on this project, the level was moved up to the logic design stage.

The system designer is responsible for translating the chip specifications into a suitable logic design using the standard cell library. The layout designer is responsible for translating that logic design into layout using the tools and ensuring that the desired performance is achieved. To achieve the performance, the layout designer may modify the schematics to implement the logic more optimally or to modify the cell designs. Modification of cell designs might in turn require a modification of the schematics. A protocol has to be agreed on which allows the layout designer to change the schematics without changing the intent of the system designer. Solution to this problem was not worked out on this project and requires further investigation.

BUMP TECHNOLOGY

INTRODUCTION

One of the fundamental limitations of high-performance VLSI-based systems is the packaging of individual devices. Systems packaging volume can be greatly reduced by making use of high-density interconnections. Typical die interconnect methods, such as wire-bonding, impose a serious limitation on operating speed due to package capacitance and the inherent self-inductance of bonding wires. The performance of CMOS-based systems can be greatly improved by reducing parasitic package lead capacitance and inductance. Typically, 30 to 50 percent of the chip power and a considerable amount of IC area are expended in large output drivers needed to overcome package parasitics.

Lead capacitance and inductance can be reduced with a direct die-attach method such as bump interconnect. With direct die-attach, connection parasitic capacitance below 0.5 picofarad (pF) is typically achieved. In conventional packaging techniques, 5 to 10 pF connection capacitance is common. Moreover, the lead self inductance (30 nanohenries (nH) per inch) is also significantly reduced, because the length of the interconnect between adjacent

chip drivers and receivers can be made very short. Low inductance is the dominant critical parameter where total interconnect lengths exceed 0.25 inch at frequencies of 50 to 100 megahertz. Because current drive requirements for inter-chip signals are reduced by bump interconnect, substantial power savings are possible.

Bump interconnect technology uses small bumps of metal or solder deposited on the die I/O pads. The die is then bonded directly to mating pads on a substrate. Bump technology can achieve interconnect densities of 2-mil centers, assuming staggered rows of bumps that are 1 mil in diameter. Conventional wire-bonding techniques require 4-mil by 4-mil pads on 8-mil centers. Typical wire-bond interconnection also limits the total number of I/O signals to the number of bonding pads that can be arranged around the perimeter of a die. Bumps, however, can be positioned anywhere on the surface of a die to dramatically improve the I/O density.

Several pairs of MOSIS TinyChip devices designed originally for other project efforts were mated using an indium bump process. These device pairs were returned to ISI for evaluation and physical inspection by sectioning and electron microscopy.

To support the next phase of bump development with silicon chips on polyimide substrates, a special test chip was designed specifically for the polyimide test structure being produced alongside the BBN substrate described in this report. This test device contained a VCO frequency source driving two experiments. The first experiment is for power dissipation analysis where the VCO drives four power inverters, with a common enable, each capable of driving 30 milliamps of current. The second experiment uses I/O drivers with varying load-handling capability to help characterize the drive requirements for the bump technology. The drive capability of the four drivers is scaled to deliver 30 milliamps, 15 milliamps, 7.5 milliamps, and 3.75 milliamps to separate outputs of the die.

A TDR experiment was also included in the test die. A line originating outside the die receives the TDR pulse. The line enters the die and extends about 1000 microns inside the die, exiting through another bump. A line on the substrate connects the line back into the die, where it then extends 2000 microns before leaving via another bump pad. This pattern is continued once more to provide 4000 additional microns of line length before exiting the die to a termination resistor on the substrate.

BUMP DIE-ATTACH EXPERIMENT

An experiment involving hybridizing the die directly to a polyimide substrate was completed. The substrate was submitted as a test coupon on the same fabrication run as the Monarch SCM substrate. Two TinyChip sites, provided on the test coupon, were connected to identical evaluation structures:

- 50-Ohm terminated input lines. These are connected by different length wires to bonding pads, allowing TDR measurements of reflections and losses caused by copper wiring on polyimide, bump junctions, and metal wiring on silicon.

- A control input for a voltage-controlled oscillator (VCO). The VCO output is buffered and presented to:
  - Identical output pads driving different capacitive loads, allowing connection inductance to be evaluated.
  - Output driver pads differing in size by factors of two, allowing connection capacitance to be measured by observing effects on rise- and fall times.

PROBE STATION ENVIRONMENT

INTRODUCTION

APT procured a low-cost automated probe station for high-speed testing and evaluation of bare dice and packages. This probe station has been greatly enhanced with an APT-developed closed-loop vision system to support automated rotation and alignment of dice and MCMs. Further enhancements include an IR laser and target recognition software to support low-cost laser customization of chips based on a Lincoln Laboratories developed linking scheme. The testing of individual bare dice in support of the Encore CDE was also completed. Use of the probe station for thermography and chemical vapor deposition is also discussed.

LASER LINKING

One of the most significant technology developments enabling the use of probe stations for system prototyping is the low-cost laser. APT purchased a laser cutting system from Alessi that mounts on the microscope camera port of a probe station. This laser, costing around $30K, replaces ultrasonic needles used for cutting metal lines on VLSI wafers. The laser has a spot size of approximately 2 microns, a cycle time of about 1 second, and sufficient power to cut metal traces through wafer passivation layers.

This instrument also has potential uses in the area of "programmable packaging." This application involves the stockpiling of standard low-cost wafers or wafer sections with generic interconnect structures, perhaps in addition to some active circuitry. These wafer-scale "packages" could be customized very quickly with the laser system by cutting traces and by connecting traces using custom links developed at Lincoln Labs, or a modification thereof. These laser-customized packages would then have custom or commercial application-specific dice mounted directly on the silicon surface.

The goal of this task was to provide a low-cost, fast-turnaround prototyping capability compatible with the high-performance laser linking facility at Lincoln Labs.

Experiments with the infrared laser for fusing links were performed with moderate success. Insights into the potential problems were gained. The major performance differences between the low-cost Nd:YAG infrared laser being used at ISI and the Lincoln Labs Argon

laser are the ISI laser's larger spot size, shorter pulse, and lack of precise control of the energy at objective - infrared tends to penetrate deeper than the green laser, requiring more control to prevent penetration into the substrate. The short pulse duration does not allow the diffusion to flow across the gap in one burst; several bursts are required to make a good link. Low power settings of the laser's power supply control do not operate the laser reliably, so energy to the link has to be controlled by de-focusing the beam and operating at higher, more stable power settings. After several attempts, a defocused beam at near maximum power setting, fired three times per link, produced links of less than 100 ohms with reasonable consistency.

TinyChip runs from several vendors were tested. Variations in link design, orientation, and application were evaluated. The designs were intended to check delay versus loading experiments, fusing of simple programmable logic modules, and direct-link resistance measurements. The logic links introduced a delay of only 100 picoseconds in a 3-micron inverter test circuit with a fan-out of two. The resistance measurement results were not as consistent. Three series of blind tests using a prescribed fusing procedure resulted in usable links 80%, 56%, and 75% of the time. These tests were conducted over several weeks, and variations in the laser or optics may were a factor. One problem appeared to be the glass overlay, which apparently varies in thickness and composition between vendors as well as within a die.

Reliable linking with this laser can be accomplished, but general application of the laser to cell re-structuring, where thousands of cells are involved, must be seriously evaluated.

Laser cutting experiments were conducted on silicon chips fabricated for the SPUR project at Berkeley. A layout error was bypassed by strategically disconnecting metal traces on the chip. Laser "edits" were performed on SPUR chips to "program" around the problem.

Additional laser experiments were performed in cooperation with University of California at Santa Barbara. These experiments were aimed at discovering the applicability of using the infrared laser to provide metal cuts in GaAs chips and to measure the effects of laser blasts on diodes and FET structures in GaAs.

LASER MEASUREMENTS

The operational characteristics of the Alessi laser and optics were suspect since the early experiments with fusing the Lincoln Labs links. The results of laser linking at ISI were very different than similar tests at LL. The required procedures for linking the LL structures were more cumbersome and less reliable at ISI. To better understand these differences and evaluate the potential for making links, a series of measurements was performed on the Alessi laser and optics to confirm its operational characteristics. Of particular interest was the power density and spot size of the laser beam at the die surface as a function of power setting for magnifications of 50X and 25X of the microscope.

The measurements were made using a broadband pyroelectric detector , op-amp and oscilliscope. The spot size was measured using two methods: by passing a focused beam through

calibrated apertures (pinholes) and by passing the beam across a partially occluded, knife edge detector.

The precision pinhole technique utilizes discs with known apertures between the laser and detector while measuring the intensity. The laser is focused on the pinhole and the detector is placed below the disc. The disc is mounted on a probe fixture that is adjustable in X, Y and Z. The fixture is adjusted for the laser to focus on the plane of the disc. Measurements are taken as the hole size is reduced, noting changes in intensity. For each intensity measurement, the disc is varied slightly in X,Y and Z seeking maximum intensity. The resulting measurements were as expected with the exception of a loss of energy when using the 50X and 25X object lenses. That is,when the expected spot size was in the order of 30 microns, the intensity readings were off by 30 percent when comparing a 200 micron pinhole with an infinite hole. The pyroelectric detector is one millimeter in diameter; therefore the loss of energy has to be due to reflections or beam diversion outside the 200 micron region. This phenomenon was not noted when using the 8X and 2.5X object lenses (Table C) and again is attributable to optical aberration(s) in the higher magnification lenses.

Another method for measuring a laser spot size is moving the detector and knife edge across the spot in precise increments while measuring the intensity at each point from. If the beam is Gaussian, the data can be fitted to an analytical approximation to erf(x), thereby extracting the beam width. Typically the beam width is defined at the half power points. However,

Power setting	Aperture size microns	Magnification	Intersity peak mv
600	∞	8X	2,400
	50		2,400
	25		1,700
	10		500
700	∞	8X	2,400
	50		2,400
	25		1,400
	10		500
600	∞	50X	500*
	50		300
	25		300
	10		200

* 50% Normal filter insulated to protect detectors

Table C. Typical Row Data from Aperture Measurements

another accepted definition is the Gaussian beam radius, where the intensity has decreased to 1/e**2 or 86..5% of the encircled power. The data with the 50X and 25X objective lenses, however, is a poor fit to the Gaussian curve (**Figure 28**), and the beam diameter estimates are not valid. Similar measurements with the 2.5X and 8X did fit the Gausssian, which again leads to the conclusion of optical aberration(s) when using the higher magnification lenses.

The results of these measurements were: maximum power .27 mj, pulse width 100 microsec., and spot size variations from 10 to 30 microns with power setting from 600 to 900 with a 50X magnification. The knife edge measurements demonstrated that the laser beam was symmetrical but non-Gaussian. Therefore the results of the spot size measurements with the apertures were not confirmed. Furthermore, both measuring techniques suggested optical aberrations, which softened the above conclusions. These conclusions suggest the laser was not operating as per specification in subtle and difficult to identify anomalies.

APT considered using a frequency doubler to modify the laser in order improve reliability of linking. The maximum laser power of .27 mj (2.7 watts at 100 microsecs.) is per specification but is insufficient for the using a non-linear crystal doubler especially when 30 percent of the beam energy is scattered outside the nominal beam. While there has been some success in laser-linking, the limitations of marginal power and poor spot shape make it difficult to preprogram a reliable procedure. The ad-hoc method of "making it work" remains the only recourse.

Despite the observed optical aberrations, the laser continues to be useful for cutting aluminum on die.

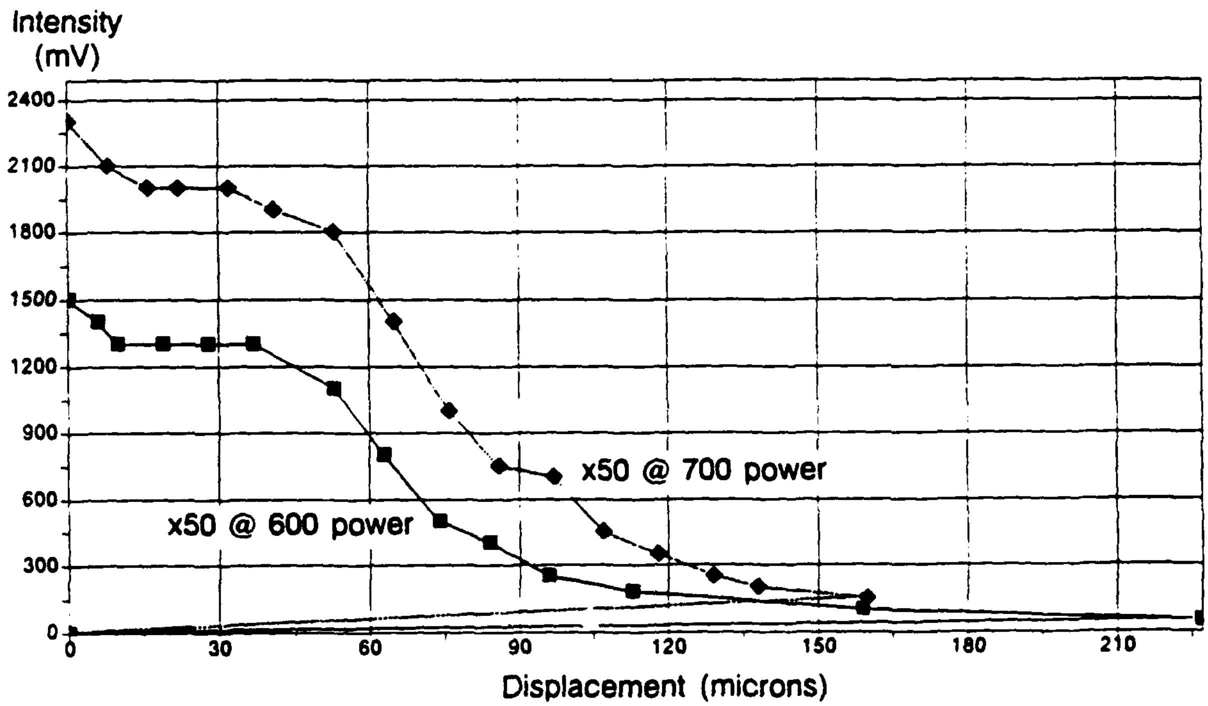

Figure 28: Typical knife-edge data

LASER RESTRUCTURING

An accurate assessment of the Alessi laser's operational performance is critical to our evaluation its role in future applications. We completed a series of measurements that confirmed our suspicions that the laser/optics were not operating as specified. The beam's energy and spot size at the objective lenses of 50X and 25X are unpredictable due to optical aberration. The net effect is a loss of 30 percent of the energy and a spot size variance (non-gaussian) from 6 to 50 microns. Data from these experiments explain the observed variance in operational results. The Alessi laser continues to be useful for cutting, but it has limited application for future laser needs of the project.

There are two approaches to upgrading the laser capability of the probe station both of which would allow it to be used as a general microsurgery device for packaging. The existing Alessi YAG laser could be upgraded through the use of a non-linear crystal to double its frequency. However, the YAG's very narrow 70 microsecond pulse width limits its applications. A more general solution involves the purchase of an Argon CW laser (488nm at 300mJ) with a shutter. This laser would cost approximately $30,000 and could support both CVD and LL links applications. With either of the above options, the optical path should be upgraded by replacing the triocular head and beam splitter and realigning the optics. This upgrade, which would cost $1,400 and double the power at the objective and would improve the accuracy of the spot size. These alternatives were identified.

VISION SYSTEM

The goal of this task is to use an image-recognition system to improve inherent positioning accuracy of a low-cost probe station for a variety of VLSI applications in design, production and testing. The vision system was developed initially to support laser linking experiments but has evolved to provide general support for other application areas.

A SUN-based vision system was successfully demonstrated. The Berkeley software for RADON transformation of image data was integrated into the ISI-written SUN application management software. Frame-digitizer hardware in the SUN was used to implement the vision system for closed-loop control of the Signatone probe station. This system corrects +/- 4-micron positional errors to within one micron, the maximum required application accuracy. The system was successfully tested with a variety of targets including links designed by ISI and Lincoln Labs as well as fiduciary marks on dice.

Project staff presented a paper, "A Vision Recognition System for High-Accuracy Position Control for Laser Reconfigurable Integrated Circuits," by W.B. Baringer, R.W. Brodersen, L. Gallenson, R. Parker, and B. White, at the twenty-second annual IEEE Asilomar Conference on Signals, Systems, and Computers, on October 31, 1988. Remote rotation control and axis mis-alignment detection has been added to the system. The RADON transformation software was modified to include projections of +/- 10 degrees in 0.2-degree steps. These projections are taken over a line segment on the die that has been manually aligned

within 10 degrees. Results of these projections are used on a "best fit" basis to instruct the rotational control on the probe station to correct the axis mis-alignment.

To improve the overall response time of the system, control character strings passed from the HP controller to the probe station were monitored and analyzed. The intent of these experiments is to provide the probe station control data directly from the SUN workstation during time-sensitive modes of operation. This approach would bypass the relatively slow operation of the HP controller and would improve system response.

As an exercise in discovering the extent of the compatibility of the low-cost laser approach to system prototyping, several ULM (Universal Logic Module) chips were procured from Lincoln Laboratories. These chips contain laser-reconfigurable logic cells that can be "programmed" under laser control to implement many different logic functions. The ULM chip is a tightly packed array of link diodes surrounded by a grid of metal-1 and metal-2 lines that potentially require cutting. Investigation of Lincoln Labs' ULM chips has produced mixed results. There are four targets (links) in each area of interest, rather than one target, as in present designs, and a new strategy is required to visually resolve these targets. The link targets were successfully handled on a 3-micron feature size chip but not on a 2-micron chip. New image-processing approaches to increase the capability for very high-density images were developed.

It is assumed that the targets or fusible links was oriented in one of two orthogonal directions and that rotation will not be necessary once the wafer or substrate has been aligned. The vision system may be required to align the wafer automatically, but there is no stringent time requirement for such setup operations.

While it is intended that the vision system need only locate a non-variant feature in the microscope field of view, there is an additional interesting concept that should be explored. The problem of testing multi-die substrates poses a challenge to a vision system. Assuming that the individual dice are mounted in known locations on a substrate but that the mounting process allows skew and slight linear misalignment, a vision system might "learn" a pattern from the manually aligned first die position and then help locate and position successive sites to support the exact alignment needed by test probe fixtures. Work at UC Berkeley in the Vision Laboratory was applied directly to this problem.

VISION SYSTEM OPERATION

The positional accuracy of the probe station is sufficient to guarantee that a target can be placed within the field of view of a television camera at the required power setting on the microscope. The field of view at 50X magnification is approximately 130 microns. The vision system, driven from a table of X-Y positions, digitizes a frame of video data and locates a particular edge-defined object within the frame. It is assumed that rotational alignment has been completed, therefore objects was oriented in one of two orthogonal directions. The vision system generates a positional correction factor in microns and commands the probe station to center the target. While it is intended that the vision system need only

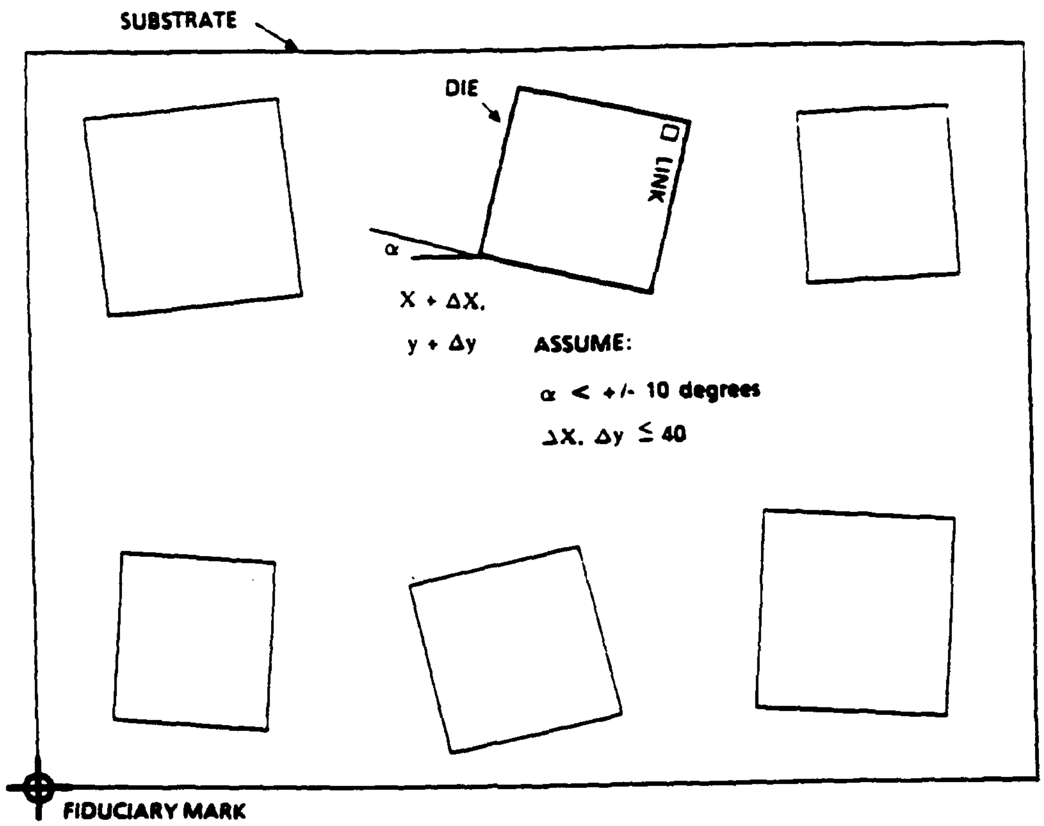

Figure 29: Multi-die Substrate Scenario

locate a non-variant feature in the microscope field of view, there is an additional interesting concept that has been explored--testing or programming multi-die substrates. The exact requirements of the application have not been defined. We have created a likely scenario to focus the vision research and to explore a broad range of requirements. The scenario was discussed in detail in previous reports. Briefly, it to say it includes learning of targets, rotational correction to within 0.1 degrees for maximum initial errors of 10 degrees, and x-y positional correction to within 1 micron. These corrections employ a Radon transform algorithm.

MULTI-DIE SUBSTRATE SCENARIO

Assuming that individual dice are mounted in known locations on a substrate, with a mounting process that allows slight linear and angular mis-alignment, the vision system can "learn" a pattern from a manually-aligned initial position. The system can then automatically locate and position the prober at successive sites in the exact alignment needed by test probe fixtures. The operation can correct for rotational as well as placement inaccuracies. (see **Figure 29**.)

The above scenario requires that the wafer or substrate be placed on the probe station and a fiduciary mark be centered in the field of view of the camera. It is assumed that dice are

positioned on the wafer with an accuracy of +/- 30 microns and the orientation is within +/- 10 degrees as compared to the wafer. The wafer need not be aligned with the coordinate system of the probe station.The program execution initializes the required devices and files and corrects the rotational alignment of the substrate. Once the substrate is properly aligned, the individual dice can be located using a file of absolute X-Y locations generated from a CIF database. Die alignment is performed automatically as programming or testing proceeds.

ROTATION ALGORITHM

The rotational algorithm is a simple extension of the two-orthogonal-projection approach used during link location.

The vision program acquires an image via the frame-digitizer and selects the region of interest (ROI), a rectangular area known to contain the target, assuming maximum positioning errors. The orientation of the die is determined by taking a series of Radon projections of the image and calculating best projection for straight lines. To minimize the number of required projections, the task is divided into three series. Projections are performed every two degrees over the ROI, followed by one-degree projection intervals to bracket the resultant value from the first series. Finally, projections are performed every 0.2 of a degree, bracketing the resultant value of the second series.

The program then instructs the probe station to rotate the wafer chuck the calculated amount and, based on trigonometric calculations, repositions the prober to the initial X-Y point on the die. The prober is then positioned to the first target point and tested for accuracy. The test is done by capturing an image of the new ROI and generating a histogram of gray-level intensities over the ROI. These histograms are correlated with learned histograms, and the positional accuracy is determined. If the probe station is in error by more than 1 micron, it is corrected by the calculated delta X and Y. Rotation now complete, this sequence is repeated using only two projections until the operation on a die site is complete.

RADON TRANSFORM

The RADON transform was selected by UC Berkeley because it is amenable to hardware implementations and can reduce the computational complexity of real-time image analysis. This transform is used to analyze lines and edges by computing projections through the image along lines at various angles. Calculations are further simplified by thresholding, which converts the gray level ROI to a black and white image, allowing 1 bit for intensity for each pixel. The projections are correlated with previously prepared projections generated by the "learning" feature of the software. These calculations are performed in the pixel domain and the error is converted from pixels to microns. If an error greater than 0.5 pixels occurs in either direction, a re-positioning command is sent to the probe station.

To conveniently operate with multiple target structures, an option for designating correlation files was added to the Vision system. The system uses a default file, XYoffplate (offset plus

template), which simplifies commands when working with a single target. The Vision system also has a command line -F option to indicate files other than the XYoffplate default.

PATTERN LEARNING

Pattern "learning" is performed with the vision system by capturing images of structures of interest, creating vector templates, editing the data to properly locate the target, and saving the data in a file for future use. This option generalizes the use of the vision system so that a variety of VLSI structures can be accurately recognized and positioned by the probe station. Learning is critical to the performance accuracy of the vision software.

The learning process begins with digitizing and processing an image of the target. A region of interest (ROI) is carefully selected to include the target structure and a border of several pixels. The digitized image of the ROI is written into a template file. Histograms are generated from several projections of the template file to ascertain the center of the target area to within one pixel. The template file plus calculated center point becomes the pattern which is compared to images of a "search area".

An option was added to make the search area size variable, which has the effect of providing a "variable-angle lens" for searching for the ROI. The need for the changeable field of view became apparent when searching for the fiduciary point during initial positioning of the die. The initial mark location can exceed specified positional tolerances if the axes of the wafer or substrate have not been precisely aligned.

SYSTEM IMPLEMENTATION

The SUN-based vision system was successfully completed during the last quarter of 1988. The UC Berkeley software for RADON transformation of image data was integrated into the ISI-written SUN application management software. A major concern with the system was its slow speed, attributed to the HP computer controller. Experiments were conducted utilizing the SUN to directly control the prober for moving the chuck in X and Y. The system was initialized using the normal HP controller and then control was switched the the SUN (RS 232 lines). The modified system was able to find and position successive links in 2.4 seconds in contrast to the 8.25 seconds required by the original system. Approximately two seconds was utilized by the RADON algorithm. This suggests our goal of zapping one target per second, the maximum capability of the laser, is within reach given the hardware implementation of the RADON algorithm. The task of completely replacing the HP controller with a SUN is significantly more complicated than the minimal functions implemented for this test. An estimate for completely eliminating the HP controller is 2-3 man-month of programming. This task is considered low priority, pending the demand for automatic targeting.

TARGET CORRELATION

The algorithm uses a cross-correlation function in identifying the proper position for the target. Experiments were conducted in utilizing this number to validate the proper position prior to a laser action. Unfortunately, the correlation number is a function of many vari-

ables, including light source, thickness of passivation, nature of target and auto-threshold level for the image. More effort is required to normalize the resulting number so that a go/nogo threshold can be set prior to laser operation.

FIDUCIARY MARKS

The selection of images for fiduciary marks is critical for accuracy and speed. Marks on dice for initializing the command sequence, as well as the image of the structure for targeting cuts and links, are currently being investigated to optimize operations. A series of tests were performed to select and recommend the appropriate structure for the vision functions. The standard plus sign (+) found on all MOSIS die ,is ideal for initial rotational correction and starting position identifier. The initiating task is made more reliable by not placing any structures within 30 microns of the plus sign. More effort is required to consider approaches for very dense arrays of links or cuts, especially with smaller features size (e.g., 2 micron).

THEMOGRAPHY

Thermal analysis of chips, substrates, and interconnect is essential to the characterization and evaluation of packaging approaches. ISI evaluated two approaches to providing thermal imaging capabilities on the probe station. The two fundamentally different approaches for thermography found in the current literature are an infra-red (IR) system and a fluorescent system. The IR system typically consists of a special CCD camera (mercury cadmium telluride detectors), optics, filters, and computer imaging hardware and software. These systems are commercially available for $50,000 to $75,000. Noise is the main limiting factor for thermal and spatial resolution in an IR system. Cooling the camera with liquid nitrogen, averaging across multiple images, and computer filtering are all used to cope with the noise problem. Each of these fixes produces some side effects on presentation, spatial resolution, or operation. The dynamic range is limited, but with initial calibration within the range of interest, measurements can be made from 0 to 1500 degrees C with about 1 percent thermal resolution and 10 micron spatial resolution. At this time there is insufficient information to estimate the complexity (or the plausibility) of interfacing components of an IR system to the prober/vision system.

A new approach to thermography is fluorescent imaging of surface temperature profiles using europium thenoyltrifluoroacetonate (EuTTA). EuTTA exhibits fluorescence that decreases with temperature when it is exposed to long-wave UV centered at 345 nm. EuTTA is spun onto the die as a polymer film and is illuminated with UV light. This light source is conveniently provided by an Hg arc lamp. The thickness of the polymer film, which is controlled by the amount of EuTTA placed on the die and the speed of rotation, is critical to the resulting presentation. The UV excitation produces a narrow-band orange output at 612 nm, whose intensity is a negative function of temperature. The visible image could be captured and processed by the ISI vision system. The intensity, and therefore the pixel values, of the image are a function of temperature. The hot image is normalized to a room-temperature image, which removes all optical anomalies and leaves a high-resolution thermal image.

While this technology is capable of 0.1 degree C resolution, it would be limited by the resolution and sensitivity of the existing ISI vision system. This technology has been developed at Bell Labs in Murray Hill, New Jersey, for obtaining temperature profiles of dice. Discussion between ISI and Bell Labs is underway to determine possible collaborative roles for extending the thermal imaging system for use in packaging applications.

CHEMICAL VAPOR DEPOSITION

Chemical Vapor Deposition (CVD) is an additive process of metals on existing substrates and is therefore a general technique for configuring dice, wafers, or substrates. It can be used to perform "microsurgery" on a variety of surfaces with great accuracy, depending on the energy source. With a laser capable of metal deposition and removal, any die or wafer can be restructured. Interest within the DARPA community in using CVD for reconfiguration has encouraged us to evaluate the application of the probe station and the vision system to CVD.

The CVD process consists of heating a surface in the presence of a metallic gas to the required temperature for deposition. The surface temperature determines the writing speed, and the beam size of the source of energy determines the line parameters. The required temperature can vary from 400 to 1000 degrees C, depending on the material to be deposited and the gas compound. A widely used deposition gas is tungsten hexafluoride (WF6) for CMOS technology. When exposed to an Argon laser (488mm) of approximately 100mw, interconnects of 1 micron thick and 8 microns wide having conductivity of 25 micro-ohm/cm can be made at writing speeds of 100 microns per second. The line width is a function of beam width and can be as narrow as 1 micron. A variety of deposition materials can be used, including Chromium, Aluminum, Molybdenum, Tantalum, Thallium, Tin, Cobalt, Silicon Oxide, Silicon Nitrate, Iodine, and commercial Diamond. The application is the prime consideration for the selection of the deposition material, but consideration must be given to the handling of the chemicals being used.

The probe station could be fitted with a vacuum chamber suitable for CVD processing. The vision system could be used to automate alignment and for pattern matching. The motor-driven stage has sufficient resolution and speed to accommodate line "writing" under computer control directly from a design database.

INTERACTIVE DIE PROBER VS. LAYOUT EDITOR TRACKING SOFTWARE

A MAGIC / prober interface has been developed which allows concurrent viewing and tracking of probe station video and die geometry images on the same workstation screen. The MAGIC / prober interface has been completed and a demonstration given.

This report describes the software developed for interfacing a die prober station to the layout editor Magic. First the goals and approach are discussed. In the next section technical details are provided. In the third section a brief users guide is included.

OVERVIEW

The goal of this project was to develop a means of comparing the actual die with the corresponding layout of a chip. The purpose of this is two fold: a) If defects are found in the die by a prober, then the user should be able to find the location of the defects in the layout and thereby analyze the effects of the defect on circuit performance; b) If the user wishes to view a specific portion of the layout in the die prober then he should be able to do so by specifying the location in the layout editor, in this case Magic. One way to achieve this goal is to provide a mechanism for communicating between the layout editor and the software that controls viewing of the die on the prober station.

Since it is expensive to maintain a die prober station, a second goal was to allow remote access via the arpanet to the die prober station located at ISI. Thus we needed a communication mechanism that will allow the layout editor running on a particular host on the arpanet to communicate with the die prober station at ISI.

APPROACH

The die prober station is linked to a SUN workstation which runs a software package to aid in viewing of the die on the workstation screen. This software initializes the position of the die being viewed by the video camera attached to the prober station. The user can then interactively enter co-ordinates of the desired area to be viewed.The software then automatically sends signals (via the vme bus) to the prober to move the prober table. The corresponding portion of the die is then displayed on the workstation screen.

To provide a link between the layout editor display and the die display on the workstation, the desired scenario is as follows: The user displays the die in one window while displaying the layout in another window. Whenever the user defines an area of interest in either window the display in the other window should move to the same location. To achieve this objective it was decided to investigate the possibility of creating a link between the layout editor program and the die prober display software while maintaining these as independent programs.

The layout editor provides interactive commands by which the coordinates of the layout area displayed can be obtained. These coordinates can be manually entered in the die prober software which sends signals to the prober station to physically move the table such that the die area displayed on the workstation has the same coordinates. Therefore, a mechanism already existed to manually achieve the above objective. As part of this project, software modules were developed for the prober control software and the layout editor to allow this communication to be done automatically.

ALTERNATIVES CONSIDERED

First a simple technique was developed to transmit the co-ordinates from the die prober to Magic. In Magic, a mechanism is available by which commands can be read from a file

instead of the keyboard. The prober software was modified to create a file with the coordinates of the die area of interest along with appropriate commands for Magic to center the layout display with these coordinates. Magic is made to read this command file and execute the commands.

This procedure was successful, however, it was cumbersome since the user has to coordinate between the prober and Magic and manually enter a command to Magic to read the file created by the prober software. Hence efforts were made to develop a more transparent link using the UNIX inter-process communication library as described in the next section.

A second problem encountered was that the co-ordinate systems of the die prober and Magic are different: the die prober works with absolute co-ordinates (microns) whereas the layout editor works with symbolic co-ordinates (lambda). A conversion routine was built in to the prober software to convert between the two coordinate systems.

A third problem was the alignment of the two coordinate systems. In Magic the origin (0,0) can occur anywhere, on the other hand the coordinate system in the die prober is always initialized such that the origin is at the lower left corner of the die. A procedure was developed so that the actual coordinates of the lower left corner of the layout can be entered (manually) in the prober software and used as an offset in the conversion routine. With the IPC interface described in the next section this procedure has also been automated.

In the final implementation other problems were encountered which required a change in the X window interface in Magic. These are described in the next section.

TECHNICAL DESCRIPTION

This section describes the techniques used in the software development as well as the individual modules. The programming effort was spent on two areas:

a) Development of interprocess communication packages (p2m and m2p) for exchanging data between the prober software and Magic. This part required an understanding of the IPC library called "sockets" in UNIX.

b) Modifications to Magic routines to allow use of interprocess communication. This part was the most time consuming due to lack of a programmers guide for Magic. However, this task proved feasible due to the excellent documentation within the Magic routines.

In this section we first give a brief introduction to the socket library. Next the packages p2m and m2p are described. Then the modifications made to Magic are described and the limitations of the software are discussed.

INTER-PROCESS COMMUNICATION LIBRARY

The UNIX operating system on the SUN workstations provides a library of functions called sockets which allow two independent processes to communicate with each other. The communication takes place as if each software module is plugged into a software socket allowing

data transmission from one module to another. The library contains three important functions: 1) to create a socket, 2) to write data to a socket and 3) to read data from a socket. The socket link is created at run time and allows each software module to be compiled and run independently.

When using the IPC library one of the software modules acts as the master and creates the socket. (The socket is physically created as a non-readable ascii file). Other software modules act as clients and "connect" to the socket. Once the socket has been created by the master and connections made by the clients, bidirectional data transfer can take place between the software modules. In this project it was decide to make Magic the master module. Client modules were developed to provide the link with the prober software. These modules are named m2p (or magic to prober) and p2m (for prober to magic) respectively (see **Figure** 30). First we describe these two client modules and then the modifications made to Magic.

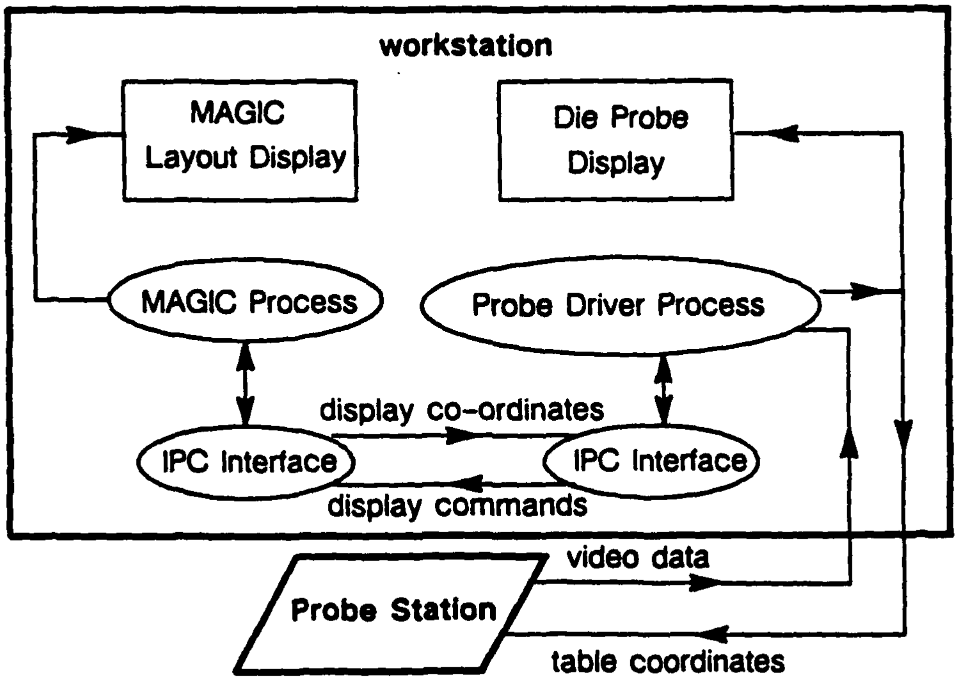

Figure 30: MAGIC / PROBER Software Perspective

The decision to make Magic the master was primarily so that Magic can be maintained as an independent program without requiring any compile time links to the prober software. The socket creation routines built into Magic are generic and can be exploited for applications other than the prober. The m2p and p2m programs are linked to the prober control routines at compile time.

Before going into details of the modules it is essential to understand the scenario in which the interprocess communication takes place: At the start of a session the user has to run Magic in one window on the workstation, which in turn sets up the socket. Next p2m or m2p

can be run in another window. With both programs running concurrently, if either wries data to the socket, the other automatically reads it. In other words if p2m or m2p write commands to the socket, magic will read these and execute them. In this sense the socket acts as yet another input device to magic (just like the mouse and the keyboard).

SOCKET USAGE BY P2M

The p2m program serves the purpose of communicating commands to Magic to force the layout display to track the die display. To do this, p2m obtains the current probe station table co-ordinates and sends Magic the *commands: box <co-ordinates> findbox and zoom* to the socket. Magic reads the commands from the socket and executes them.

The box command in Magic defines the co-ordinates of the layout area to be displayed, findbox centers the layout display on this box and zoom causes the boxed area to fill as much of the window as possible with the given aspect ratio.

In order to keep this communication transparent to the user, every time the user moves the die display (by moving the prober table), the p2m program automatically carries out the above tasks. Note that Magic and p2m run concurrently and Magic continuously polls the socket for inputs just as it polls the keyboard. Thus the user does not have to take any action to transmit the commands to Magic or for Magic to read these commands. The entire procedure is executed automatically.

While p2m is independent of Magic and only assumes the creation of a socket by Magic at run time, p2m is linked to the prober control software at compile time. Thus execution of p2m also executes the prober control software. Since this is custom software it was deemed appropriate to make this link. If necessary one make p2m independent of the prober software also and set up a three way link amongst Magic, p2m and the prober controller. However this would have taken more time and no advantages could be found with this approach.

Note that p2m only requires a one way communication between the prober and Magic. The other module (m2p) requires a two way communication as described below.

SOCKET USAGE BY THE M2P PACKAGE

The m2p program causes the die display to track the Magic layout display. Since updating the die display is a relatively slow process, it was decided not to update the die display every time the layout display is changed by the user. Instead, m2p provides an interactive interface by which the user specifically executes a command instructing m2p to update the die display. When this command is executed, m2p obtains the layout box co-ordinates from magic and provides these to the prober control program.

The way m2p achieves this is to send a box command to Magic through the socket. In response Magic executes the box command which causes it to write the coordinates of the current box location (in the layout display) to the socket. These are read from the socket by m2p and passed to the prober display controller.

Note that the socket is bidirectional so the command from m2p to Magic and the box co-ordinates from Magic to m2p are communicated through the same socket.

To exploit the IPC library it was necessary to edit several routines in Magic as described below. These changes will become part of the Magic version released with the LagerIV system by Berkeley.

MODIFICATIONS TO MAGIC

In order to exploit the socket IPC library, several routines in Magic had to be modified.

The normal mechanism for executing commands in Magic is to either type them in on the keyboard or read them from a file. For the IPC facility it was necessary to introduce a third mechanism, namely that of reading commands from a socket. Furthermore it is desirable that Magic reads the commands in an asynchronous manner just as it does from the keyboard, i.e., whenever commands are written to the socket by a remote procedure (p2m or m2p) Magic should read them and execute them.

To implement this facility we had to analyze the command entry and execution structure in Magic. It was found that the file grXinput.c in Magic has several routines that handle all input devices for Magic. In this file the GrXWInitialize() function initializes Magic to accept inputs from the keyboard and the mouse. This routine has been modified so it creates a socket with the IPC library and allows

Magic to accept data coming in on the socket from a remote process. The p2m and m2p programs connect to this socket at run time.

The GrXinput() function in the same file was modified so it polls the socket in addition to the keyboard and the mouse for user inputs. Thus any data transmitted to the socket by a remote process (m2p and p2m in this case) is read by Magic as if it had been entered on the keyboard. When a command is detected on the socket further handling of the command is done identically as with keyboard entered commands so no further modifications were necessary.

One problem encountered was that when Magic receives the commands on the socket from a remote process (such as p2m) it cannot uniquely decide which window to execute the command in even if the cursor is in the layout window.

It appears that when commands are entered from the keyboard, in addition to interpreting and executing the command magic also updates an internal window pointer based on the cursor location. This action did not take place when commands were entered through the socket. It was unclear what would be a good general solution to this problem. Accordingly Brian Richards added a few functions in the X window interface of Magic, which keep track of the last window in which a command was executed and set a default pointer to this window. When commands from the socket are executed Magic uses this default pointer. This strategy works successfully under the assumption that in this application the user only

has one layout window open and the cursor is in that window. A more flexible approach can be implemented in a future extension of the project.

The above modifications in Magic allow remote processes to send commands to Magic via the socket. For the m2p program it is also desirable to have a communication the other way: when the box command is sent by m2p to the socket Magic prints the co-ordinates of the box in the layout window to the console using a set of text I/O routines. It is also necessary to send this data to the socket for m2p. Two alternatives were considered. One was to modify the text I/O routines so all messages sent to the console are also sent to the socket. This was not considered appropriate since in most cases the socket would get flooded with data that is not required and would have to be flushed out. The alternative approach that was implemented was to modify the I/O routine corresponding only to the box command. If data printed by other commands is required to be transmitted to the socket in the future, the same modification can be made to the appropriate command I/O routine.

LIMITATIONS

The above modifications to Magic actually allow any remote process to communicate with Magic just as p2m and m2p do by connecting to the socket. However, currently m2p and p2m are two separate programs and in one session only one of them can be connected to the socket created by Magic. For most applications this is not a draw back since either the layout has to track the die prober or vice-versa. However if simultaneous tracking is required both ways p2m and m2p can be merged into one program.

A second limitation is that all programs have to be executed on the same machine for the socket IPC to work. Thus both Magic and the probe controller software must run on the same machine. The probe controller software has to run on the machine that is interfaced to the prober station. Therefore users at remote sites have to rlogin to the prober workstation at ISI and execute Magic and the prober software on that machine. Since all displays (Magic as well as the die prober) are X window based the user at a remote site can view the displays on his machine. The disadvantage is that the workstation controlling the prober station has more load on it since it also has to run Magic.

The socket library also allows sockets to be created for communication between processes on two different hosts on the internet. In a future enhancement the socket functions in Magic, m2p and p2m can be modified to exploit this capability if it is desirable for the remote user to be able to run magic locally.

BRIEF USER GUIDE

The m2p and p2m programs allow designers to visually compare the chip layout against an image of the die obtained with a wafer probe station as described below:

a) p2m: This package can send commands to Magic to center the layout display on the same area as being viewed by the probe station.

b) m2p: This package can send commands to the prober to view the area enclosed by the box in the Magic display.

To use these programs the user executes Magic and displays the desired layout. Next the p2m or m2p program is executed depending on the objective. (Magic has to be executed first since it creates the socket interface).

The p2m package is fully automatic. When run initially it obtains layout coordinates from Magic to compute conversion factors between Magic co-ordinates and the probe station co-ordinates. Thereafter whenever the probe station table is moved, the probe image on the workstation is updated and simultaneously commands are sent to Magic to shift the layout display accordingly. No user intervention is necessary.

The m2p package is interactive since the user may not always wish the probe station table to be moved when the layout display is moved. When an update of the probe image is desired, the user executes a command in m2p. The layout display co-ordinates are automatically obtained from Magic by m2p and in turn sent to the probe station to move the table so that the die image on the workstation tracks the layout display.

DIE TESTING

One of the most severe problem areas in advanced packaging is the testing of bare dice. Many advanced packaging techniques intended to maximize performance require that dice not be packaged in conventional single-die packages. Foundries are generally willing to provide bare individual dice but are unwilling to provide wafers. The problem of testing these commercial dice and any custom prototype dice produced, for example, through MOSIS, becomes a problem of testing individual bare dice. Custom probe cards designed for unique die pad locations can be purchased relatively inexpensively for manual and automatic probe stations. These probe cards and probe stations are designed to support wafer testing of dice but testing individual dice is very difficult. The difficulty lies in handling, aligning, and holding the individual die. During this reporting period, APT has demonstrated an approach to automatically testing individual bare dice using the low-cost probe station environment. This demonstration served two purposes. First, APT had a specific die testing problem in support of the Encore CDE, and second, the demonstration showed feasibility for a remote testing capability in the critical area of individual bare dice.

One approach to packaging the memory module required by the Encore CDE involved the use of SRAM die that were available only in individual die form. These dice had to be functionally screened to maximize the post processing yield before being assembled into an MCM. APT procured a 28-probe, custom probe card designed to test the SRAM dice (Cypress CY7C192) on the APT automated probe station. A mechanical jig was designed and fabricated to roughly align the individual dice on the probe station chuck. An interface cable was designed and fabricated to connect the probe card to an IMS functional tester. Two tests were written to provide a checkerboard pattern and an address test (mod 16). The rotation

and alignment software developed by APT for the probe station was used to provide precise control of die position for proper test probe alignment.

While the components of the individual dice testing process were automated, the actual demonstration required operator intervention at several points. The automated control of the probe station lowers the chuck and moves it forward toward the user to facilitate loading and unloading of the test die. While a robot with a vacuum arm could be used to move the die from a waffle-pack to the outstretched chuck, this capability is not included in the present system. In this demonstration the operator moves the die from the waffle-pack to the alignment fixture mounted to the chuck. The operator then initiates the automated load, align, and test functions successively. After the tests are complete, the operator initiates the unload command and the chuck presents the die for removal. All of these commands can be initiated remotely and could be fully automated if a capability for moving the die to and from the waffle-pack were added.

An automated individual die testing service would require the development of a software control environment to replace the operator intervention required in the completed demonstration. A simple robot with the ability to pick up dice from a waffle-pack and move laterally a fixed distance to the probe station chuck would also have to be added. An x-y table would hold the waffle-pack and position it under the robot vacuum arm to align the vacuum arm with the proper die location. The unique feature that allows these low-cost components to work for this application is the automatic alignment software and closed-loop probe station control developed by APT.

Assuming that the automated testing service described above were implemented, the cost to "tool" a new die type would be limited to the cost of developing another probe card. Probe cards for devices with 28 pins cost about $400.00, or about $14 per pad. While this cost is reasonable, it would make good economic sense to establish standard frames for large pin count devices to minimize the re-tooling costs.

INDIVIDUAL DIE THINNING

Commercial chip design houses invest significant effort in reducing the area of chips. Since die yield is inversely proportionally to die size, die shrinks are a common method to reduce costs as chip designs mature. A system designer wishing to minimize the overall volume of a system is limited by the combined areas of the individual die in the system. However, the designer can reduce system volume by thinning these die. Thinning is done commonly on uncut wafers and commercial companies exist that provide wafer lapping services. It is often the case, however, that devices are only available as individual, cut bare die. This significantly complicates the thinning process.

A technology experiment was performed at ISI in exploring a low-cost approach to thinning individual bare die. SRAM die were thinned to progressively smaller thicknesses and functionally tested at each step of the process.

A relatively low cost 64K x 4 SRAM die was selected for this experiment because SRAMs are easy to test and easy to interface to test equipment. This experiment included selecting a commercial SRAM die and attempting thinning to .015", .010", .005", .002". As expected, the thinnest die warped significantly because of the stresses in the over-glass. SRAM die approximately .005" thick remained flat and passed functional testing.

TAB DEVELOPMENT

PROTOTYPE TAB

TAB packaging technology is well suited to high pin count, higher performance chips. It supports testing and is cheap to produce in high volume. It is, however, very expensive to tool for prototype and low volume applications. Commercial systems houses report that it takes one year and costs about $100K to provide tooling for a new high volume TAB package.

A new technology experiment was undertaken at ISI to investigate a low cost prototype TAB manufacturing capability. This effort was coupled to the thin die experiment described elsewhere in this report. A TAB design was completed for the SRAM die used for the thinning experiments. This TAB design was fabricated and used to package the thinned SRAM die.

The TAB parts were used to provide information on thinning process yield. The intent of this experiment was not only to validate the TAB process but also to demonstrate a low profile packaging approach compatible with the thin die.

ANALOG CORRELATOR IC

As a vehicle for evaluating high-frequency analog packaging, a CMOS analog correlator chip architected by Dr. Asad Abidi of UCLA for spread spectrum decoding was designed at ISI. This device uses CMOS analog technology provided by Dr. Abidi and Ramon Gomez. The differential circuits employed provide noise margin that allows the correlator to be used as a component of a digital CMOS chip. The correlator was fabricated using the MOSIS SCEA (double-metal, double-poly) 2-micron analog process.

In this device, the analog input was clocked through a 15-stage switched-capacitor tapped delay line. Tap weights were set by pass-transistor switches. The switch outputs were differential and were fed to a switched-capacitor summing tree. The final output of the summing tree was the correlation signal, which was buffered for output off the chip.

A differential op-amp design was designed and simulated. The HSPICE simulations of hand-extracted geometry show an open loop gain of 71.5 dB. The 3 dB point was at 63.1 Khz and unity gain is at 316 MHz.

Each switched-capacitor delay stage uses an op-amp and a single 29.5fF capacitor for both input and feedback, guaranteeing unity gain. During Φ_1 high, both outputs of the amp swing to the common mode voltage, 2.75V. During this time the capacitors are charged from the

differential inputs. During Φ_2 high, the capacitors were switched into the feedback loop. Each delay line stage also contains a reference generator for all bias levels used in the delay segment.

PN switches are made from two 16 switch elements sections. Only switches that correspond to slave stages of the delay line were used. Switch control lines were driven off-chip for testing.

The summing tree is made up of switched-capacitor units, one with 15 inputs, the other with 16 inputs. These are in turn are summed by a 2 input stage. Each summing stage contains a reference generator. In each of the summing units, the clocks were exchanged between stages, causing alternating cells to function as master and slave. Unused inputs are tied to a 2.75V reference distributed throughout the chip.

The analog summing element was fully differential, but for simplicity a single-ended circuit is described. During Φ_2, the output of the op-amp was connected to the summing junction, forcing the junction to 0 volts. One side of C1 and C2, 29.5fF capacitors, are also tied to the summing junction. The other sides of C1 and C2 were connected to the input voltages Va and Vb, respectively. During Φ_1, the Vb side of C2 was grounded, forcing the summing junction to a potential of -Vb, while the Va end of C1 was switched to the output of the op-amp. To force the summing junction back to 0V, the amplifier output must rise to Va+Vb volts.

The output buffer is comprised of 2 CMOS stages set for unity gain. HSPICE simulations predict buffer performance of 42.3dB open loop gain, unity-gain at 630MHz, and a 3dB point at 6.3MHz. The phase margin was predicted to be 12.7 degrees and the slew rate predicted to be 14V/µS. The buffer was stable into a 16pF load.

TEST BOARD

A test board for the correlator chip was designed with selectable clocks. Clock signals are shifted from TTL to the +1V to +6V levels used by the correlator. The board also contains a PN generator implemented by a [4,1] feedback shift register. The serial PN output is modulated by the chip-rate clock using a D flip-flop and an XOR gate. The output from the modulator was fed to a voltage divider to set the input level of a unity gain phase splitter. The outputs of the phase splitter were fed to the inputs of the correlator chip.

TESTING

To reduce capacitive loading during operation, a laser was used to cut internal probe pads at the summing tree outputs. All summing tree connections were left intact. Testing found that 11 out of 12 chips were functional, although care is required in the power-up sequence to avoid a non-destructive latch-up condition. Cycling power correctly corrects this problem.

A second test set the master clock to 2MHz and was used to drive both the PN generator and the correlator chip. It was found that 100% correlation produces an output 1.0V above the

reference and 100% anti-correlation produces an output −1.25V below the reference. Due to a pipeline created by the depth of the clocked summing tree, the correlator output was delayed 4 clocks after the [1,1,1,1] state of the PN generator.

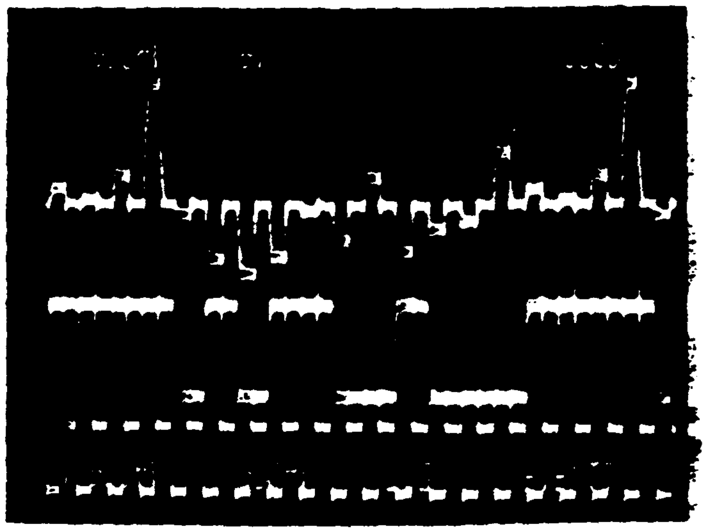

2.0MHz CORELATION 9-3-91

Figure 31. Correlator output signal (f=2MHz)

Figure 31 shows the correlator output at 2Mhz. In this picture, an oscilloscope differential input was directly connected to the correlator outputs.

5X FET 2

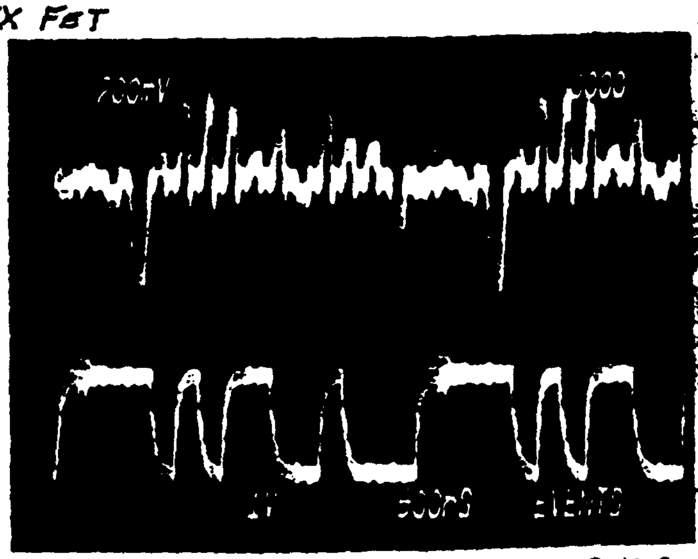

5MHz CHIP=1.6Vpp 9-12-91

Figure 32. Correlator output signal (f=5MHz)

At frequencies above 5MHz, a 5X buffered FET-input probe was used to examine the output. Study of the waveforms such as those in **Figure 32** indicates that the slew rate of the output buffer amps was 10V/μs not 14V/μs as predicted by HSPICE. From this information we infer that the internal circuitry of the correlator runs faster than the buffer amps.

RESULTS

The essential circuitry of the correlator has 63 internal op-amps and two unity gain output buffers amps. The power consumption of the correlator chip is given in the following table.

Supply	I(mA)	V	Power	freq (MHz)
Vdd	21.98	4.97	109 mW	2.0
	21.51	4.97	107 mW	5.0
	21.25	4.97	106 mW	10.0
	22.32	4.97	110 mW	15.0
NWELL	37μA	5.99	222 μW	2.0
	37μA	5.99	222 μW	5.0
	33μA	5.99	198 μW	10.0
	22μA	5.99	132 μW	15.0

CONCLUSIONS

The test results show that, operating at a peak frequency of 15MHz, total device power was on the order of 110 milliwatts. The 63 op amps on the chip occupy an active area 2715μ by 2639μ, and at 1.5 mW per op amp represent one of the most aggressively scaled high speed, precision analog building blocks demonstrated to date.

Significant loss of signal amplitude was observed in internal stages apparently caused by amplifier sensitivity to loading.

LOANER PROGRAMS

APT has established the Tester Loaner Program and, more recently, the CAD Loaner Program to provide universities with easy access to low-cost commercial test equipment and CAD tools, and to encourage universities to include issues of testing in VLSI education.

This program is administered on a voluntary basis by the Integrated Systems Laboratory at ISI. The purpose of the Loaner Program is to provide VLSI designers with access to low-cost functional test systems in support of both education and research. The Loaner Program also provides a forum for discussion of all aspects of testing through a column published in the MOSIS Users Group Newsletter (MUG).

TESTER LOANER PROGRAM

A limited number of loaner test systems are available on a first-come-first-served basis to non-commercial MOSIS users. These systems are provided by commercial test system man-

ufacturers who have joined the Loaner Program. The Loaner Program also gives users an opportunity to purchase systems directly from the manufacturer at a substantial discount. To date we have circulated 6 testers to about 20 universities. Eighty percent of the users that try the lowest-price tester take advantage of the special pricing and purchase the unit.

TEST HARDWARE

Under an agreement with CADIC, several loaner Model 4100 "state testers" were provided for evaluation and rotated through Pennsylvania State University, the University of Tennessee, the University of Southern California, the University of Utah, the University of California at Berkeley, the University of California at San Diego, the University of California at Los Angeles, the University of Pennsylvania, and Washington University (where it is presently being evaluated). The Integrated Measurement Systems loaner has been evaluated by the University of Tennessee, Syracuse University, the University of California at San Diego, and the University of Washington. Dartmouth University and the Oregon Graduate Center are waiting in the rotation queue.

CAD LOANER PROGRAM

The CAD Loaner Program is similar to the Tester Loaner Program. The terms of the CAD Loaner program are negotiated separately with each CAD supplier, but the basic terms are as follows: commercial CAD suppliers that are affiliated with the CAD Loaner Program supply CAD tools to universities free of charge for a period of one year. At the end of the one-year period, the university may return the software without obligation or may purchase the software at an 80 percent discount from list price. During the loaner period (and after purchase, if the university exercises its option), no annual maintenance fee is charged, but support is limited to occasional phone consultation. A training class is to be held twice a year at the CAD vendor's location. The class is free to CAD Loaner universities, but travel to the class must be provided by the university.

The basic CAD Loaner Program has been established, and agreements were signed with two CAD suppliers. ViewLogic is offering its schematic front-end and simulation package, and Task Technologies is offering its PCB/hybrid routing package.

The CAD Loaner Program has supplied multiple copies of Sun- and PC-based ViewLogic systems to the University of Southern California, the University of California at Los Angeles, the Jet Propulsion Laboratory, and the University of California at Santa Barbara. Since the recent announcement that ViewLogic has been selected by MOSIS as the front-end capture and simulation package for their netlist-to-parts service, the level of CAD Loaner inquiries has dramatically increased.

PUBLICATIONS AND PRESENTATIONS

PUBLICATIONS

"High Density Systems Modules (HDSM), An MCM-based approach to building high-performance multiprocessors"
Proceedings of the IEEE/NSF 1991 MCM Workshop
March 28-29, 1991, Santa Cruz, CA.

Packaging White Paper.
June 1988. Unpublished. Referenced in
"Rapid Prototyping Facilities in the U.S. Manufacturing Research Community"
1990, Report of the Manufacturing Studies Board, National Research Council.

"Advanced Production Technologies Project: Semi-Annual Technical Report "
May 1988; November 1988; April 1989; November 1989; March 1990; October 1990; March 1991; October 1991.

"Where Do We Go From Here: Emerging Programs in Support of Education"
Proceedings VLSI Conference and Exposition, Summer 1989.

"A Vision Recognition System for High Resolution Position Control for Laser Reconfigurable Integrated Circuits"
Baringer W.B., and Broderson, R.W., University of California, Berkeley; Gallenson L., Parker R.H. and White B., Information Sciences Institute, University of Southern California. Invited paper for the 22nd Annual IEEE Asilomar Conference on Signals, Systems, and Computers. October 31 - November 2, 1988.

"Functional Testing"
Published in the MOSIS Users' Group Newsletter re-published in IEEE Circuits and Devices magazine, May 1988; Volume 4; Number 3.

"Testers and Testing Issues"
MOSIS Users' Group Newsletter, #2, April 1988, #3, May 1988.

Technical Report for the Steering Committee,
Packaging Subcommittee of the Tera-op Technology Working Group,
March 31 - April 1, 1988, Boston, Massachusetts.

Report on the National Security Industrial Agency
Taskforce on Packaging and Interconnect Meeting.
February 25-26 1988, Anaheim, California.

Report of the Test Working Group,
NSF Workshop on Undergraduate VLSI Education,
Robert Parker, Chairman. November 30 - December 1, 1987.

DARPA Semi-Annual Technical Report,
KITSERV and System Engineering, November 1987.

PRESENTATIONS

"Packaging Alternatives for Embedded Variants"
Presented at the Honeywell Supercomputing Workshop, December 19, 1991

"Progress in Advanced Production Technologies"
Presented at DARPA VLSI Contractors Meeting, November 15, 1991

"Multiprocessor Packaging Alternatives"
Presented at Intel Supercomputers, November 1, 1991

"High Density Systems Modules (HDSM), An MCM-Based Approach to Building High-Performance Multiprocessors"
Presented at the IEEE/NSF 1991 MCM Workshop, March 28-29, 1991

"Advanced Production Technologies briefing at Encore Computer"
Presented at Encore Computer, November 30, 1990

"Advanced Production Technologies briefing at SCC"
Presented at Space Computer Corporation, November 15, 1990

"Packaging Techniques for Heterogeneous Systems"
Presented at the DARPA VLSI Contractors Meeting, October 3-5, 1990

"Suggested Improvements to SEM-E & JIAWG Programs Implemented with Multi-chip Modules (MCM)"
Presented at Sandia National Labs, September 4, 1990

"Packaging Technology Access"
Presented at the DARPA ISAT Meeting, August 3, 1990

"MCM Technology"
Presented at the 1990 Microelectronic System Education Conference & Exposition August 1, 1990

"Advanced Production Technology Briefing"
Presented to the Space Computer Corporation, June, 19, 1990

"Testing Issues for Educators"
Presented to the VLSI Educational Conference and Exposition, August 24, 1988